arXiv:1703.06323v1 [math.NA] 18 Mar 2017

ROBUST AND SCALABLE DOMAIN DECOMPOSITION SOLVERS FOR UNFITTED FINITE ELEMENT METHODS

SANTIAGO BADIA† AND FRANCESC VERDUGO‡

Abstract. Unfitted finite element methods have a great potential for large scale simulations, since avoid the generation of body-fitted meshes and the use of graph partitioning techniques, two main bottlenecks for problems with non-trivial geometries. However, the linear systems that arise from these discretizations can be much more ill-conditioned, due to the so-called small cut cell problem. The state-of-the-art approach is to rely on sparse direct methods, which have quadratic complexity and thus, are not well-suited for large scale simulations. In order to solve this situation, in this work we investigate the use of domain decomposition preconditioners (balancing domain decomposition by constraints) for unfitted methods. We observe that a straightforward application of these preconditioners to the unfitted case has a very poor behavior. As a result, we propose an enhancement of the classical BDDC methods based on 1) a modified (stiffness) weighting operator and 2) an improved definition of the coarse degrees of freedom in the definition of the preconditioner . These changes lead to a robust and algorithmically scalable solver able to deal with unfitted grids. A complete set of complex 3D numerical experiments show the good performance of the proposed preconditioners.

Keywords: Unfitted finite elements, embedded boundary methods, linear solvers, parallel computing, domain decomposition

Contents

1. Introduction 2
2. Unfitted FE method 4
2.1. Cell partition and subdomain partition 4
2.2. Model problem and space discretization 5
2.3. Numerical integration in cut cells 7
3. Domain decomposition for body-fitted FE meshes 8
3.1. Sub-assembled problem 8
3.2. Coarse DOFs 8
3.3. Transfer operator 9
3.4. Preconditioner and condition number bounds 9
4. Domain decomposition solvers for unfitted FE meshes 11
4.1. Breakdown example 11
4.2. Enlarged coarse spaces and adaptive BDDC methods 12
4.3. Stifness weighting operator 12
4.4. New coarse objects: variant 1 12
4.5. New coarse objects: variant 2 13
5. Numerical experiments 14
5.1. Setup 14

Date: March 21, 2017.

† Universitat Politècnica de Catalunya, Jordi Girona1-3, Edifici C1, E-08034 Barcelona & Centre Internacional de Mètodes Numèrics en Enginyeria, Parc Mediterrani de la Tecnologia, Esteve Terrades 5, E-08860 Castelldefels, Spain E-mail: sbadia@cimne.upc.edu. SB gratefully acknowledges the support received from the Catalan Government through the ICREA Acadèmia Research Program.

‡ Universitat Politècnica de Catalunya, Jordi Girona1-3, Edifici C1, E-08034 Barcelona & Centre Internacional de Mètodes Numèrics en Enginyeria, Parc Mediterrani de la Tecnologia, Esteve Terrades 5, E-08860 Castelldefels, Spain E-mail: fverdugo@cimne.upc.edu.

5.2. "Popcorn flake" example. 17
5.3. Multi-body example: sphere, popcorn flake, block, array of blocks and spiral. 20
6. Conclusions 22
References 23

1. Introduction

The use of unfitted finite element methods (FEMs) is an appealing approach for different reasons. They are interesting in 1) coupled problems that involve interfaces (e.g., fluid-structure interaction [8] or free surface flows), or 2) situations in which one wants to avoid the generation of body-fitted meshes. These type of techniques have been named in different ways. Unfitted FEMs for capturing interfaces are usually denoted as eXtended FEM (XFEM) [11], whereas these techniques are usually denoted as embedded (or immersed) boundary methods, when the motivation is to simulate a problem using a (usually simple Cartesian) background mesh. Recently, different realizations of the method have been coined in different ways, depending on the way the numerical integration is performed, how Dirichlet boundary conditions are enforced on no-matching surfaces, or the type of stabilization, if any, being used. To mention some examples, the finite cell method combines an XFEM-type functional space, a numerical integration based on adaptive cell refinement and full sub-cell integration, to perform integration on cut cells, and a Nitsche weak imposition of boundary conditions [38]. CutFEM makes also use of Nitsche's method, but includes additional "ghost penalty" stabilization terms to improve the the stability of the algorithm [15]. The huge success of isogeometrical analysis methods (spline-based discretization) and the severe limitations of this approach in complex 3D geometries will probably increase the interest of unfitted methods in the near future [27].

Another justified motivation to turn our attention to unfitted methods is related to the simulation of large scale complex problems. Nowadays, we have at our disposal excellent (non)linear solver libraries, that can efficiently exploit the largest supercomputers today in implicit simulations of problems governed by partial differential equations (PDEs). In addition, many algorithmic ideas have been proposed to extend this scalability to the forthcoming exascale supercomputers [5, 22, 30, 30, 40]. In this situation, the bottleneck of the simulation pipeline (at least for industrial applications with non-trivial geometries) is being shifted to the body-fitted mesh generation step and the unstructured mesh partition (usually based on graph partitioning algorithms), since it is hard to extract concurrency from the underlying algorithms. In fact, the state-of-the-art graph partitioning algorithms and implementations being used in computational science and engineering (CSE) have not changed much in the last 20 years, since the groundbreaking work by Karypis and Kumar [29]. The decline of such approaches is in contraposition with the increasing interest on octree meshes, space-filling curves, and parallel adaptive FEMs [2, 9, 18, 26]. Using "forests of octrees", one can define a coarse mesh that topologically describes the original geometry but still keep to a great extent the excellent scalability in the generation and partition of octree meshes (as soon as the coarse mesh is small enough). In any case, it involves coarse body-fitted hexahedral meshes (even harder to generate than tetrahedral meshes), algorithms to put nodes on the surface (complicating for high order geometrical approximation), and it does not scale with the size of the coarse mesh (usually, the coarse mesh is replicated in all processors [18]). Fortunately, recent research has been devoted to solve these issues, e.g., using octree-type libraries for tetrahedral meshes [17] or untangling techniques in high order mesh generation [46]. Another approach to solve these issues is to consider unfitted FEMs on background octree meshes, which is the approach advocated in this work.

The use of unfitted methods on background octree Cartesian meshes avoids the need to define body-fitted meshes, and can exploit efficient and scalable space-filling curve algorithms for the partition and load re-balancing steps. The surface detection by marching cubes-type algorithm and edge intersection (e.g., using level set methods or distant functions) can easily be performed in parallel. Thus, the meshing and partition steps can certainly benefit from the use of unfitted meshes, eliminating two

bottlenecks in CSE simulations. In turn, such schemes complicate the numerical integration, imposition of Dirichlet boundary conditions, and the linear solver phase. Integration on cut elements complicates the implementation of the FE code, even though different alternatives have been proposed in the literature to address this problem (see, e.g., [44]). Probably, the main issue is the definition of Dirichlet boundary conditions on the unfitted surfaces. Different techniques have been considered so far. In this work, we consider Nitsche's method to the weak imposition of the Dirichlet boundary conditions [24], since it provides arbitrarily high order schemes (due to consistency). The drawback is the fact that the jump penalty coefficient required to have stability must be computed at the cut element level and can blow up for cut elements (see, e.g., [14, 19]) with close to zero support (small cut cell problem). Additionally, one can consider ghost penalty stabilization on cut cells, in order to have stability in an extended domain with full elements. This way, the Nitsche's penalty coefficient and the condition number of the resulting matrix are bounded above. The price to pay is a weakly non-consistent algorithm. The method in [16] is based on jump penalty terms that require to compute high order derivatives on faces for high order finite elements (FEs), certainly complicating the implementation of the methods. Alternatively, the scheme in [14] only involves projections on the cut cells, where different local projections can be considered, e.g., relying on Scott-Zhang projector [3]. We note that the ghost penalty stabilization *is not needed* in practice for coercive problems to have a stable and convergent algorithm. The situation is unclear for saddle point problems, and a field of current research [23].

The main showstopper up to now for the succesful application of unfitted methods for realistic applications is the linear solver step. The condition number of the resulting linear system does not only depend on the characteristic size of the background mesh elements, but also on the characteristic size of the cut elements; cut elements can be arbitrarily small and can have arbitrarily high aspect ratios for small cut cell situations. Enforcing some reasonable thresholds, one can use robust sparse direct solvers for these problems. However, sparse direct methods are very expensive, due to their quadratic complexity, which is especially dramatic at large scales. Scalability is also hard to get, especially at very large scales. For large scale applications with body-fitted meshes in CSE, iterative solvers, usually Krylov solvers combined with preconditioners, is the natural way to go. Scalable preconditioners at large scales are usually based on algebraic multigrid (AMG) and domain decomposition solvers [5, 22, 30, 30, 40]. Unfortunately, the lack of robust and scalable iterative solvers for unfitted FEM has limited the applicability of unfitted methods in real applications. It is well-known that some rudimentary methods, e.g., Jacobi preconditioners, can solve the issue of the ill-conditioning that comes from cut elements. However, these preconditioners are neither scalable nor optimal. Different serial linear solvers for unfitted FEMs have been recently proposed (see, e.g., [12, 19, 25, 35]). The methods in [12, 35] consider a segregation of nodes into healthy and ill nodes, and considers a domain decomposition solver for the ill nodes. The method is only applied to 2D problems in [35] and serial computations, and the domain decomposition solver is not scalable, due to the fact that no coarse correction is proposed. The method in [12] proposes to use an AMG method for the healthy nodes. A specific-purpose serial AMG solver is designed in [25] and a serial incomplete factorization solver is considered in [19].

Motivated by the lack of robust (with respect to the element cuts) and scalable parallel solvers for unfitted methods, we develop in this work a domain decomposition solver with these desired properties based on the balancing domain decomposition by constraints (BDDC) framework. First, we provide a short description of the mathematical analysis that proves the scalability of BDDC methods on body-fitted meshes. Next, we show an example that proves that the use of standard BDDC methods cannot be robust with respect to the element cuts. In fact, the method performs dramatically bad when straighforwardly applied to unfitted meshes. Next, we propose (based on heuristic arguments and numerical experimentation, but motivated from the mathematical body of domain decomposition methods) a modified BDDC method. First, we consider a new weighting operator, relying on the diagonal of the (sub-assembled stiffness matrix). The use of this weighting operators proves to be essential for the robustness of the solver. Next, we consider enhanced coarse spaces, with a sub-partition of the edge constraints. Since no mathematical analysis is available for domain decomposition

methods with stiffness weighting, we have performed a comprehensive set of numerical experiments on 3D geometries with different levels of complexity. The methods are not only robust and algorithmically weakly scalable, but what is more surprising, the number of iterations seems to be independent of the domain shape too. Furthermore, the cost of the modified weighted operator is identical to the standard one, and the algorithm has the same building blocks as the standard BDDC preconditioner. This is good, since one can use the extremely scalable implementation of these methods, e.g., in the `FEMPAR` library [1, 6]. In any case, the cost of the new method can differ from the standard one in the number of coarse degrees of freedom (DOFs). We have observed that in practice, the CPU cost of the modified formulation is very close to the one of the standard preconditioner. In fact, as we increase the size of the coarse problem of the modified solver (which is robust for unfitted methods) tends to the one of the standard preconditioner (which performs very bad for unfitted meshes). We note that the preconditioner proposed herein could also be readily applied to XFEM-enriched interface problems and other discretization techniques (e.g., discontinuous Galerkin methods).

Let us describe the outline of this work. In Sect. 2 we show the unfitted FE method considered in this paper, in particular our choice to integrate in cut cells and the surface intersection algorithm. (In any case, the preconditioners proposed later on do not depend on these choices, and can be used in other situations). We also state the model problem, its discretization, and the mesh and subdomain partitions. The standard BDDC preconditioner and its building blocks for body-fitted meshes are presented in Sect. 3. In Sect. 4, we show a breakdown example of the preconditioner on unfitted meshes, an expensive solution that is provably robust, and a cheap solution based on the stiffness weighting and (optionally) two slightly larger coarse spaces. A complete set of complex 3D numerical experiments are included in Sect. 5 to show experimentally the good performance of the proposed preconditioner. Finally, some conclusions are drawn and future work is described in Sect. 6.

2. Unfitted FE method

2.1. **Cell partition and subdomain partition.** Let $\Omega \subset \mathbb{R}^d$ be an open bounded domain, with $d \in \{2,3\}$ the number of spatial dimensions. For the sake of simplicity and without loss of generality, we consider that the domain boundary is defined as the zero level-set of a given scalar function ϕ^{ls}, namely $\partial\Omega \doteq \{x \in \mathbb{R}^d : \phi^{\text{ls}}(x) = 0\}$. We note that the problem geometry could be described using 3D CAD data instead of level-set functions, by providing techniques to compute the intersection between cell edges and surfaces (see, e.g., [34]). In any case, the way the geometry is handled does not affect the domain decomposition solver presented below. Like in any other immersed boundary method, we build the computational mesh by introducing an *artificial* domain Ω_{art} such that 1) it has a simple geometry easy to mesh using Cartesian grids and 2) it includes the *physical* domain $\Omega \subset \Omega_{\text{art}}$ (see Fig. 1a).

Let us construct a partition of Ω_{art} into *cells*, represented by θ_{art}. We are interested in θ_{art} being a Cartesian mesh into hexahedra for $d = 3$ or quadrilaterals for $d = 2$, even though unstructured background meshes can also be considered. Cells in θ_{art} can be classified as follows: a cell $e \in \theta_{\text{art}}$ such that $e \subset \Omega$ is an *internal cell*; if $e \cap \Omega = \emptyset$, e is an *external cell*; otherwise, e is a *cut cell* (see Fig. 1b). The set of interior (resp. external and cut) cells is represented with θ_{in} and its union $\Omega_{\text{in}} \subset \Omega$ (resp. $(\theta_{\text{ext}}, \Omega_{\text{ext}})$ and $(\theta_{\text{cut}}, \Omega_{\text{cut}})$). Furthermore, we define the set of *active cells* as $\theta_{\text{act}} \doteq \theta_{\text{in}} \cup \theta_{\text{cut}}$ and its union Ω_{act}. Let us also consider a subdomain partition Θ of Ω_{act}, obtained by *aggregation of active cells* in θ_{act}, i.e., there is an element partition $\theta^{\omega}_{\text{act}} \doteq \{e \in \theta_{\text{act}} : e \subset \omega\}$ for any $\omega \in \Theta$. The interface of the subdomain partition is $\Gamma \doteq \cup_{\omega\in\Theta}\partial\omega \setminus \partial\Omega$.

To generate the subdomain partition in the context of immersed boundary methods is simple, since Cartesian meshes are allowed. For uniform Cartesian meshes, the elements can be easily aggregated into uniform subdomains as shown in Fig. 1d, whereas for adaptive Cartesian meshes the elements can be efficiently aggregated using space-filling curves [2]. In both cases, the partitions can be generated efficiently in parallel without using highly costly graph-based partitioning techniques [28], which are otherwise required for complex unstructured body-fitted meshes. In fact, the elimination of both body-fitted mesh generation and graph partitioning is the main motivation of this work.

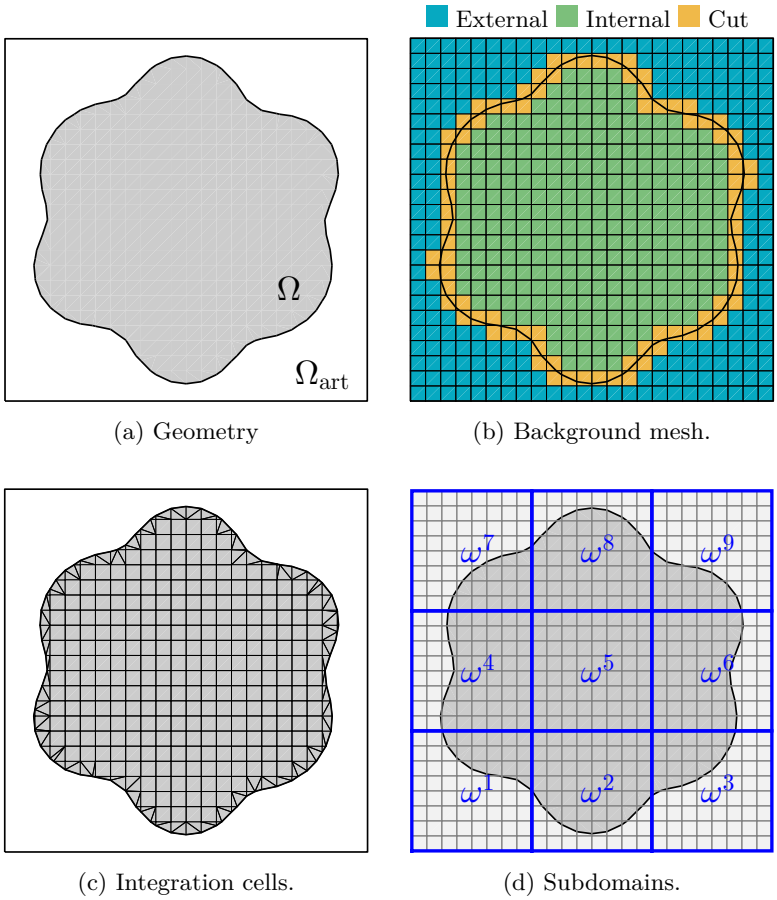

(a) Geometry (b) Background mesh.

(c) Integration cells. (d) Subdomains.

FIGURE 1. Immersed boundary setup.

2.2. **Model problem and space discretization.** Let us consider as a model problem the Poisson equation with Dirichlet and Neumann boundary conditions: find $u \in H^1(\Omega)$ such that

$$-\Delta u = f \quad \text{in} \quad \Omega, \qquad u = g^{\mathrm{D}} \quad \text{on} \quad \Gamma_{\mathrm{D}}, \qquad \nabla u \cdot n = g^{\mathrm{N}} \quad \text{on} \quad \Gamma_{\mathrm{N}}, \tag{1}$$

where $(\Gamma_{\mathrm{D}}, \Gamma_{\mathrm{N}})$ is a partition of the domain boundary (the Dirichlet and Neumann boundaries, respectively), $f \in H^{-1}(\Omega)$, $g^{\mathrm{D}} \in H^{1/2}(\Gamma_{\mathrm{D}})$, and $g^{\mathrm{N}} \in H^{-1/2}(\Gamma_{\mathrm{N}})$.

For the space discretization, we consider H^1-conforming FE spaces on conforming meshes. (The discontinuous Galerkin (DG) case will not be considered in this work, but we could also develop preconditioned iterative linear solvers for unfitted DG formulations by combining the ideas herein with the BDDC preconditioner for DG formulations in [21].) For unfitted grids, it is not possible to include Dirichlet conditions in the approximation space in a strong manner. Thus, we consider Nitsche's method [10, 37] to impose Dirichlet boundary conditions weakly on Γ_D. It provides a consistent numerical scheme with optimal converge rates (also for high-order elements) that is commonly used in the immersed boundary community [42]. We define the space $\bar{\mathbb{V}} \subset H^1(\Omega)$ as the global FE space related to the mesh θ_{act}. Further, we define the FE-wise operators:

$$\mathcal{A}_e(u,v) \doteq \int_{e\cap\Omega} \boldsymbol{\nabla} u \cdot \boldsymbol{\nabla} v \, \mathrm{d}V + \int_{\Gamma_{\mathrm{D}}\cap e} \left(\beta uv - v\,(\boldsymbol{n}\cdot\boldsymbol{\nabla} u) - u\,(\boldsymbol{n}\cdot\boldsymbol{\nabla} v)\right) \, \mathrm{d}S,$$

$$\ell_e(v) \doteq \int_{\Gamma_{\mathrm{D}}\cap e} \left(\beta v g^{\mathrm{D}} - (\boldsymbol{n}\cdot\boldsymbol{\nabla} v)\, g^{\mathrm{D}}\right) \, \mathrm{d}S,$$

defined for a generic mesh element $e \in \theta_{\text{act}}$. Vector $\boldsymbol{n}$ denotes the outwards normal to $\partial\Omega$. The bilinear form $\mathcal{A}_e(\cdot,\cdot)$ includes the usual form resulting for the integration by parts of (1) and the additional term associated with the weak imposition of Dirichlet boundary conditions with Nitsche's method. The right-hand side operator $\ell_e(\cdot)$ includes additional terms related to Nitsche's method. We will make abuse of notation, using the same symbol for a bilinear form $\mathcal{A} : \bar{\mathbb{V}} \times \bar{\mathbb{V}} \to \mathbb{R}$ and its corresponding linear operator $\mathcal{A} : \bar{\mathbb{V}} \to \bar{\mathbb{V}}'$ as $\langle \mathcal{A}u, v\rangle \doteq \mathcal{A}(u,v)$.

The coefficient $\beta > 0$ is a mesh-dependent parameter that has to be large enough to ensure the coercivity of $\mathcal{A}_e(\cdot,\cdot)$. As shown in [19], coercivity is mathematically guaranteed by considering an element-wise constant coefficient β such that

$$\beta_e \geq C_e \doteq \sup_{v\in\bar{\mathbb{V}}} \frac{\mathcal{B}_e(v,v)}{\mathcal{D}_e(v,v)}$$

for all the mesh elements $e \in \theta_{\text{cut}}$ intersecting the boundary Γ^{D}. In previous formula, β_e is the value of β restricted to element e, whereas $\mathcal{D}_e(\cdot,\cdot)$ and $\mathcal{B}_e(\cdot,\cdot)$ are the forms defined as

$$\mathcal{D}_e(u,v) \doteq \int_{e\cap\Omega} \boldsymbol{\nabla} u \cdot \boldsymbol{\nabla} v \ \mathrm{d}V, \quad \text{and} \quad \mathcal{B}_e(u,v) \doteq \int_{\Gamma_{\text{D}}\cap e} (\boldsymbol{n}\cdot\boldsymbol{\nabla} u)\,(\boldsymbol{n}\cdot\boldsymbol{\nabla} v) \ \mathrm{d}S.$$

Since $\bar{\mathbb{V}}$ is finite dimensional, and $\mathcal{D}_e(\cdot,\cdot)$ and $\mathcal{B}_e(\cdot,\cdot)$ are symmetric and bilinear forms, the value C_e (i.e., the minimum admissible coefficient β_e) can be computed numerically as $C_e \doteq \tilde{\lambda}_{\max}$, being $\tilde{\lambda}_{\max}$ the largest eigenvalue of the generalized eigenvalue problem (see [19] for details):

$$\mathcal{B}_e x = \tilde{\lambda} \mathcal{D}_e x. \tag{2}$$

In this work, we chose $\beta_e \doteq 2\tilde{\lambda}_{\max}$ in all the examples in order to be on the safe side, as suggested in [19]. Let us note that both operators have the same kernel, the space of constants. Thus, the generalized eigenvalue problem (2) is well-posed (see [19] for more details).

The global FE operator $\bar{\mathcal{A}} : \bar{\mathbb{V}} \to \bar{\mathbb{V}}'$. and right-hand side term $\ell \in \bar{\mathbb{V}}'$ are stated as the sum of the element contributions, i.e.,

$$\bar{\mathcal{A}}(u,v) \doteq \sum_{e\in\theta_{\text{act}}} \mathcal{A}_e(u,v), \quad \ell(v) \doteq \sum_{e\in\theta_{\text{act}}} \ell_e(v), \qquad \text{for } u,v \in \bar{\mathbb{V}}.$$

Further, we define $b : \bar{\mathbb{V}}' \to \bar{\mathbb{V}}$ as $b(v) \doteq f(v) + g^{\text{N}}(v) + \ell(v)$, for $v \in \bar{\mathbb{V}}$.

The global problem can be stated in operator form as: find $u \in \bar{\mathbb{V}}$ such that $\bar{\mathcal{A}}u = b$ in $\bar{\mathbb{V}}'$. After the definition of the set of DOFs $\bar{\mathcal{N}}$ and the corresponding FE basis (of shape functions) $\{\phi_a(\boldsymbol{x})\}_{a\in\bar{\mathcal{N}}}$ that span $\bar{\mathbb{V}}$, the previous problem leads to a linear system to be solved. At large scales, sparse linear systems are usually solved with Krylov iterative methods [41] due to the quadratic complexity of sparse direct methods. Here, we can consider a conjugate gradient (CG) iterative method since $\mathcal{A}$ is symmetric positive definite. The convergence rate of the conjugate gradient method is very sensitive to the condition number of the operator, namely $k_2(\bar{\mathcal{A}}) \doteq |\bar{\mathcal{A}}|_2/|\bar{\mathcal{A}}^{-1}|_2$, where $|\bar{\mathcal{A}}|_2$ denotes the operator 2-norm. As shown in [19] (under over-optimistic assumptions), the condition number scales as

$$k_2(\bar{\mathcal{A}}) \sim \min_{e\in\theta_{\text{cut}}} \eta_e^{-(2p_e+1-2/d)}, \qquad \eta_e \doteq \frac{|e\cap\Omega|}{|e|},$$

where p_e is the polynomial degree of the interpolation at element e. Thus, arbitrarily high condition numbers are expected in practice since the position of the interface cannot be controlled. In consequence, the convergence rate of the iterative solver is expected to be very slow and highly sensitive to the position of the geometry unless a robust preconditioner is considered. Our goal is to develop a BDDC preconditioner able to 1) deal with the bad conditioning properties associated with cut elements, 2) be quasi-optimal with respect to h (the characteristic element size of θ_{act}), and 3) be weakly scalable (not dependent on the number of subdomains for a fixed subdomain problem size). Attaining these objectives, we will have at our disposal a solver for embedded boundary systems that will efficient exploit distributed memory resources for large scale problems.

2.3. **Numerical integration in cut cells.** We still need to define a strategy to carry out the numerical integration in cut cells. In contrast to conventional conforming FEs, the mesh elements cannot be used as integration cells if they are cut by the boundary. In this case, the integration has to be performed only in the part of the cell $e \in \theta_{\text{cut}}$ that intersects the domains of interest, namely $e \cap \Omega$, for volume integrals, and $e \cap \partial\Omega$ for surface integrals associated with boundary conditions. In order to integrate over these geometrically complex domains, we build for each $e \in \theta_{\text{cut}}$ one tessellation (i.e. a tetrahedral/triangular mesh) of the intersected volume $e \cap \Omega$ and another one for the intersected surface $e \cap \partial\Omega$ (see Fig. 1c). Then, we consider a Gaussian quadrature rule within each tessellation cell to evaluate the integrals.

The construction of the tessellations of $e \cap \Omega$ and $e \cap \partial\Omega$ requires some further elaboration. They are defined by 1) the coordinates of the tessellation points, and 2) the connectivity of the tessellation cells. For each cut element, we define the tessellation points as the union of the element nodes and the intersection points between the element edges and the embedded boundary. Since we consider level set methods, it is straightforward to determine the intersection points on element edges. Assuming a linear variation of the level-set function within each mesh element, an edge would be intersected if and only if the level set function has different signs at the two edge endpoints. Moreover, if the edge is intersected, the location of the intersection point is computed by interpolating the values of level set function at the edge end points. Second, we efficiently built the connectivity of the tesselation cells following the spirit of the marching cubes algorithm [32, 36], where the cell connectivities are taken from a small set of pre-computed patterns. In 3D, a total of 2^8 intersection cases are possible, taking into account that an hexahedron has 8 nodes and that the value of the level-set function is either positive (including 0) or negative. Among these cases, only 14 are unique if rotations, symmetries, and uncut elements are eliminated. For each case, volume and surface tesselations are generated using a Delaunay triangulation algorithm and stored in a look up table (see Figs. 2 and 3). Using this pre-computed information, we perform a loop on the cut elements of the mesh. For each element, we determine which is the intersection case, and we extract the corresponding cell connectivities. These connectivities together with the previously defined tessellation points completely determine the tessellations within each cut element.

Note that the Delaunay algorithm is performed only once for creating the cell connectivities for each intersection pattern. This pre-computed information is used later at each cut element of the mesh without computing any Delaunay triangulation again. Yet similar, our approach is different to the classical marching cubes algorithm since 1) we generate both surface and volume tessellations (the marching cubes algorithm only provides surface meshes) and 2) we compute the exact position of the intersection points for each cut element (whereas the classical marching cubes method usually approximates the intersection location as the edge midpoint).

We chose this numerical integration approach, based on the tessellation of cut elements, for its simplicity and efficiency, but alternative methods (see, e.g., [44]) could also be used without any change in the domain decomposition solver presented below.

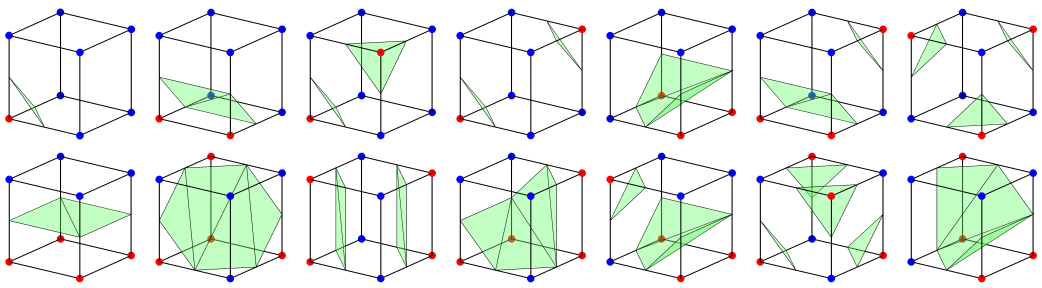

FIGURE 2. Surface tessellations for the marching cubes cases.

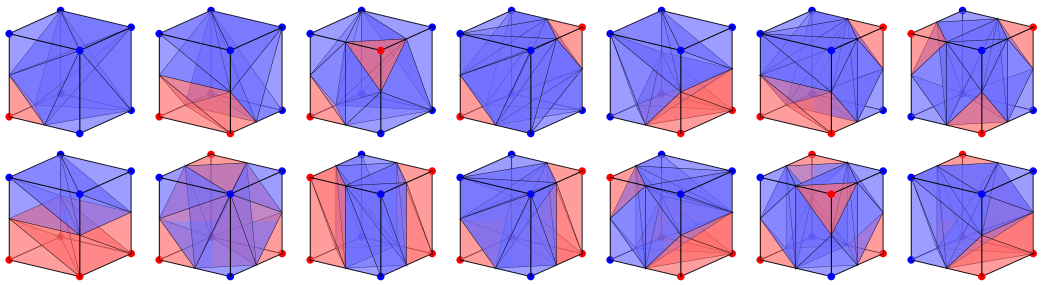

FIGURE 3. Volume tessellations for the marching cubes cases.

3. Domain decomposition for body-fitted FE meshes

In this section, we briefly introduce the BDDC method as usually considered in the literature for conventional body-fitted meshes [20]. The description of the algorithm is concise and we mainly focus on the parts that are modified later in Sect. 4 to deal with unfitted meshes. We refer the reader to [13, 33] for a detailed exposition of BDDC methods and to [4, 5] for the practical implementation details. We have also included a short proof of the condition numbers for body-fitted meshes. It is important for the subsequent discussion to explain the breakdown of the standard algorithm for unifitted meshes and motivate robust modifications. In this section, $\Omega \equiv \Omega_{\text{act}}$ and the Dirichlet boundary conditions are strongly enforced in the definition of the FE spaces.

3.1. Sub-assembled problem.

Non-overlapping domain decomposition preconditioners rely on the definition of a sub-assembled FE problem, in which contributions between subdomains have not been assembled. In order to do so, at every subdomain $\omega \in \Theta$, we consider the FE space $\mathbb{V}_\omega$ associated to the element partition $\theta^\omega_{\text{act}}$ with homogeneous Neumann boundary conditions on $\partial\omega \cap \partial\Omega$. One can define the subdomain operator $\mathcal{A}_\omega(u,v) = \sum_{e\in\theta^\omega_{\text{act}}} \mathcal{A}_e(u,v)$, for $u, v \in \mathbb{V}_\omega$.

Subdomain spaces lead to the sub-assembled space of functions $\mathbb{V} \doteq \Pi_{\omega\in\Theta}\mathbb{V}_\omega$. For any $u \in \mathbb{V}$, we define its restriction to a subdomain $\omega \in \Theta$ as u_ω. Any function $u \in \mathbb{V}$ can be represented by its unique decomposition into subdomain functions as $\{u_\omega \in V_\omega\}_{\omega\in\Theta}$. We also define the sub-assembled operator $\mathcal{A}(u,v) \doteq \Pi_{\omega\in\Theta}\mathcal{A}_\omega(u_\omega, v_\omega)$.

With these definitions, $\bar{\mathbb{V}}$ can be understood as the subspace of functions in $\mathbb{V}$ that are continuous on the interface Γ, and $\bar{\mathcal{A}}$ is the Galerkin projection of $\mathcal{A}$ onto $\bar{\mathbb{V}}$. We note that θ_{act} and the FE type defines $\bar{\mathbb{V}}$, whereas Θ is also required to define the local spaces $\{\mathbb{V}_\omega\}_{\omega\in\Theta}$ and the sub-assembled space $\mathbb{V}$, respectively. We consider Lagrangian FE spaces, where DOFs are associated to nodes (spatial points). We represent the set of nodes related to the FE space $\mathbb{V}$ with $\mathcal{N}$ (analogusly for $\mathcal{N}_\omega$).

3.2. Coarse DOFs.

A key ingredient in domain decomposition preconditioners is to classify the set of nodes of the FE space $\bar{\mathbb{V}}$. The interface ∂e of every FE in the mesh θ_{act} can be decomposed into vertices, edges, and faces. By a simple classification of these entities, based on the set of subdomains that contain them, one can also split the interface Γ into faces, edges, and vertices (at the subdomain level), that will be called geometrical objects. We represent the set of geometrical objects by Λ. In all cases, edges and faces are open sets in their corresponding dimension. By construction, faces belong to two subdomains and edges belong to more than two subdomains in three-dimensional problems. This classification of Ω into objects automatically leads to a partition of interface DOFs into DOF objects, due to the fact that every DOF in a FE does belong to only one geometrical entity. These definitions are heavily used in domain decomposition preconditioners (see, e.g., [45, p. 88]).

Next, we associate to some (or all) of these geometrical objects a *coarse* DOF. In BDDC methods, we usually take as coarse DOFs mean values on a subset of objects Λ_O. Typical choices of Λ_O are $\Lambda_O \doteq \Lambda_C$, when only corners are considered, $\Lambda_O \doteq \Lambda_C \cup \Lambda_E$, when corners and edges are considered, or $\Lambda_O \doteq \Lambda$, when corners, edges, and faces are considered. These choices lead to three common variants of the BDDC method referred as BDDC(c), BDDC(ce) and BDDC(cef), respectively. This classification

of DOFs into objects can be restricted to any subdomain $\omega \in \Theta$, leading to the set of subdomain objects $\Lambda_O(\omega)$.

With the classification of the interface nodes and the choice of the objects in Λ_O, we can define the coarse DOFs and the corresponding BDDC space. Given an object $\lambda \in \Lambda_O(\omega)$, let us define its restriction operator $\mathcal{R}_\lambda^\omega$ on a function $u \in \mathbb{V}_\omega$ as follows: $\mathcal{R}_\lambda^\omega(u)(\xi) = u(\xi)$ for a node $\xi \in \mathcal{N}_\omega$ that belongs to the geometrical object λ, and zero otherwise. We define the BDDC space $\tilde{\mathbb{V}} \subset \mathbb{V}$ as the subspace of functions $v \in \mathbb{V}$ such that the constraint

$$\alpha_\lambda^\omega(v_\omega) \doteq \int_\lambda \mathcal{R}_\lambda^\omega(v_\omega) \qquad \text{is identical for all } \omega \in \text{neigh}(\lambda), \tag{3}$$

where $\text{neigh}(\lambda)$ stands for the set of subdomains that contain the object λ. (The integral on λ is just the value at the vertex, when λ is a vertex.) Thus, every $\lambda \in \Lambda_O$ defines a coarse DOF value (3) that is continuous among subdomains. Further, we can define the BDDC operator $\tilde{\mathcal{A}}$ as the Galerkin projection of $\mathcal{A}$ onto $\tilde{\mathbb{V}}$.

3.3. **Transfer operator.** The next step is to define a transfer operator from the sub-assembled space $\mathbb{V}$ to the continuous space $\bar{\mathbb{V}}$. The transfer operator is the composition of a weighting operator and a harmonic extension operator.

(1) The weighting operator $\mathcal{W}$ takes a function $u \in \mathbb{V}$ and computes mean values on interface nodes, i.e.,
$$\mathcal{W}u(\xi) \doteq \frac{\sum_{\omega \in \text{neigh}(\xi)} u_\omega(\xi)}{|\text{neigh}(\xi)|},$$
at every node $\xi \in \mathcal{N}$ of the FE mesh θ_{act}, where $\text{neigh}(\xi)$ stands for the set of subdomains that contain the node ξ. It leads to a continuous function $\mathcal{W}u \in \bar{\mathbb{V}}$. It is clear that this operator only modifies the DOFs on the interface. We denote this weighting as *topological weighting*, since it is defined from the topology of the mesh and the subdomain partition.

(2) Next, let us define the bubble space $\mathbb{V}_0 \doteq \{v \in \mathbb{V} : v = 0 \text{ on } \Gamma\}$ and the Galerkin projection $\mathcal{A}_0$ of $\mathcal{A}$ onto $\mathbb{V}_0$. We also define the trivial injection $\mathcal{I}_0$ from $\mathbb{V}_0$ to $\bar{\mathbb{V}}$. The *harmonic* extension reads as $\mathcal{E}v \doteq (I - \mathcal{I}_0\mathcal{A}_0^{-1}\mathcal{I}_0^T\bar{\mathcal{A}})v$ for a function $v \in \bar{\mathbb{V}}$, where I is the identity operator in $\bar{\mathbb{V}}$. This operator corrects interior DOFs only. The computation of $\mathcal{A}_0^{-1}$ involves to solve local problems with homogeneous Dirichlet boundary conditions on Γ. The set of discrete global harmonic functions is $\mathbb{V}_0^\perp \doteq \{v \in \mathbb{V} : \ \mathcal{A}(v, w) = 0, \ \forall w \in \mathbb{V}_0\}$.

The transfer operator $\mathcal{Q} : \mathbb{V} \to \bar{\mathbb{V}}$ is defined as $\mathcal{Q} \doteq \mathcal{E}\mathcal{W}$.

3.4. **Preconditioner and condition number bounds.** With all these ingredients, we are now in position to define the BDDC preconditioner. This preconditioner is an additive Schwarz preconditioner (see, e.g., [45, Ch. 2]), with corrections in $\mathbb{V}_0$ and the BDDC correction in $\tilde{\mathbb{V}}$ with the transfer $\mathcal{Q}$. As a result, the BDDC preconditioner reads as:

$$\mathcal{M} \doteq \mathcal{I}_0\mathcal{A}_0^{-1}\mathcal{I}_0^T + \mathcal{Q}\tilde{\mathcal{A}}^{-1}\mathcal{Q}^T.$$

For body fitted meshes and 2nd order elliptic problems, the BDDC preconditioned operator, namely $\mathcal{M}\bar{\mathcal{A}}$, has a condition number bounded as follows (see [33]).

Theorem 3.1. *The operator $\mathcal{M}\bar{\mathcal{A}}$, where $\mathcal{M}$ can be the preconditioner BDDC(c) or BDDC(ce) in 2D and BDDC(ce) and BDDC(cef) in 3D, has a condition number bounded by:*

$$k_2(\mathcal{M}\bar{\mathcal{A}}) \leq \tilde{C}\left(1 + \log^2\left(\frac{H}{h}\right)\right), \tag{4}$$

where $\tilde{C}$ is a constant independent from the sizes h and H (the characteristic mesh and subdomain size). In the case of BDDC(ce), every face $\lambda_F \in \Lambda_F$ must have an edge $\lambda_E \in \Lambda_E$ on its boundary, i.e., $\lambda_E \subset \partial\lambda_F$.

Proof. The BDDC method can be casted in the abstract additive Schwarz setting (see [13, Ch. 7] and [45, Ch. 2]). As a result, the condition number of the preconditioned operator is $k_2(\mathcal{M}\bar{\mathcal{A}}) = \frac{\lambda_{\max}}{\lambda_{\min}}$, with $\lambda_{\min} = 1$ and

$$\lambda_{\max} \doteq \max_{\tilde{v}\in\tilde{\mathbb{V}}} \frac{\langle \mathcal{A}\mathcal{Q}\tilde{v}, \mathcal{Q}\tilde{v}\rangle}{\langle \mathcal{A}\tilde{v}, \tilde{v}\rangle} \leq 1 + \max_{\tilde{v}\in\tilde{\mathbb{V}}\cap\mathbb{V}_0^\perp} \frac{\langle \mathcal{A}(\mathcal{Q}\tilde{v} - \tilde{v}), \mathcal{Q}\tilde{v} - \tilde{v}\rangle}{\langle \mathcal{A}\tilde{v}, \tilde{v}\rangle},$$

where we have used the triangle inequality and the energy minimization property of discrete harmonic functions in the last bound (see [33]). In order to bound $\lambda_{\max}$, we need some further elaboration. We can decompose any discrete harmonic function $v \in \mathbb{V}_0^\perp$ as the sum of object functions $v = \sum_{\lambda\in\Lambda}\sum_{\omega\in\text{neigh}(\lambda)} \mathcal{E}\mathcal{R}_\lambda^\omega(v)$. For body-fitted meshes, it is clear that

$$\mathcal{R}_\lambda^\omega(v) = 0, \quad \text{on} \quad \partial\omega \setminus \lambda, \quad \forall\lambda \in \Lambda, \quad \forall\omega \in \text{neigh}(\lambda). \tag{5}$$

Further, given $v \in \mathbb{V}_0^\perp$, $\lambda \in \Lambda$, and subdomains $\omega, \omega' \in \text{neigh}(\lambda)$, we define the jump function $\mathcal{J}_{\omega,\omega'}^\lambda(v) \in \mathbb{V}_0^\perp$ that is equal to zero in all subdomains but ω, where it take the value $\mathcal{J}_{\omega,\omega'}^\lambda(v)(\xi) \doteq v_\omega(\xi) - v_{\omega'}(\xi)$ if $\xi \in \lambda$ and zero otherwise. Thus, we get

$$w \doteq \mathcal{Q}v - v = \sum_{\lambda\in\Lambda}\sum_{\omega\in\text{neigh}(\lambda)} \mathcal{E}\mathcal{R}_\lambda^\omega(\mathcal{Q}v - v) = \sum_{\lambda\in\Lambda}\sum_{\omega\in\text{neigh}(\lambda)}\sum_{\omega' \text{neigh}(\lambda)} \frac{1}{|\text{neigh}(\lambda)|}\mathcal{J}_{\omega,\omega'}^\lambda(v). \tag{6}$$

Let us consider a constant γ such that

$$\|\mathcal{J}_{\omega,\omega'}^\lambda(v)\|_{\mathcal{A}} \leq \gamma\|v\|_{\mathcal{A}}, \quad \text{for any} \quad v \in \tilde{\mathbb{V}},\ \lambda \in \Lambda,\ \omega, \omega' \in \text{neigh}(\lambda). \tag{7}$$

The maximum eigenvalue can readily be bounded by $\lambda_{\max} \leq C(1+\gamma)$ using the triangle inequality an the fact that $|\text{neigh}(\lambda)|$ is uniformly bounded. The remaining ingredient is to find a γ satisfying (7) for all objects in Λ. We observe that w vanishes on corners for $v \in \tilde{\mathbb{V}}$, since corner values are identical on all subdomains. Thus, only edge and face objects must be considered. Let us consider first edge terms. We consider $\lambda \in \Lambda_E$ and $\omega, \omega' \in \text{neigh}(\lambda)$. We note that the mean value of any $v \in \tilde{\mathbb{V}}$ is continuous between subdomains for BDDC(ce) and BDDC(cef), i.e., $\alpha_\lambda^\omega(v) = \alpha_\lambda^{\omega'}(v)$. We have:

$$\begin{aligned}\|\mathcal{J}_{\omega,\omega'}^\lambda(v)\|_{\mathcal{A}} &= |\mathcal{J}_{\omega,\omega'}^\lambda(v)|_{H^1(\omega)} = C|\mathcal{J}_{\omega,\omega'}^\lambda(v)|_{H^{\frac{1}{2}}(\partial\omega)} \\ &\leq |\mathcal{R}_\lambda^\omega(v) - \alpha_\lambda^\omega(v)|_{H^{\frac{1}{2}}(\partial\omega)} + |\mathcal{R}_\lambda^{\omega'}(v) - \alpha_\lambda^{\omega'}(v)|_{H^{\frac{1}{2}}(\partial\omega')}.\end{aligned}$$

We have used the trace theorem in $H^1(\omega)$, the fact that $\mathcal{J}_{\omega,\omega'}^\lambda(v)$ is discrete harmonic, $\mathcal{J}_{\omega,\omega'}^\lambda(v) = 0$ on $\partial\omega \setminus \lambda$, and the fact that the mean value of functions in $\tilde{\mathbb{V}}$ are continuous between subdomains. We proceed analogously for faces in BDDC(cef). Finally, using the fact that

$$|\mathcal{R}_\lambda^\omega(v) - \alpha_\lambda^\omega(v)|^2_{H^{\frac{1}{2}}(\partial\omega)} \lesssim \left(1 + \log^\rho\left(\frac{H}{h}\right)\right)|v|^2_{H^1(\omega)},$$

where $\rho = 1$ for edges and 2 for faces, we bound the edge terms for BDDC(ce) and edge/face terms for BDDC(cef) (see [45, Lemma 4.26]). The bound for faces in BDDC(ce) differs from the fact that $\alpha_\lambda^\omega(v) \neq \alpha_\lambda^{\omega'}(v)$ in general. It can be bounded similarly relying on the assumption that there exists an edge in $\partial\lambda$ that is in Λ_O (see [45, p. 182] for more details). □

In a weak scaling test (i.e., refining the mesh and the partition at the same rate), H/h is constant and so is the condition number bound regardless of the problem size. This property leads to an optimal weak scaling of the preconditioner (i.e., to increase h and H in the same proportion) when used within a conjugate gradient iteration. That is, the number of iterations needed to solve the problem up to a certain (relative) tolerance is proportional to the condition number $k_2(\mathcal{M}\bar{\mathcal{A}})$, and as the condition number is bounded by a constant independent of the problem size, so are also the number of iterations. The coarse DOFs can be reduced, including only a subset of edges in Λ_O. We refer to [31] for more details and the extension to linear elasticity problems.

4. Domain decomposition solvers for unfitted FE meshes

The BDDC method presented in the previous section can be applied verbatim to unfitted meshes. In that case, the computational mesh is θ_{act} and the extended domain Ω_{act}. It obviously requires specific strategies for integrating the problem matrices within cut elements and to impose Dirichlet conditions on unfitted boundaries, but the BDDC solver can be applied unaltered. The question is, however, whether the performance of BDDC is strongly deteriorated or not when moving from body-fitted to unfitted meshes.

The bound for the condition number in (4) cannot be extended to unfitted methods. One of the problems is that (5) for the physical subdomain $\omega \cap \Omega$, namely

$$\mathcal{R}_\lambda^\omega(v) = 0, \quad \text{on} \quad \partial(\omega \cap \Omega) \setminus \lambda, \quad \forall \lambda \in \Lambda, \quad \forall \omega \in \text{neigh}(\lambda),$$

is not true for unfitted meshes, since a partition can have cut cells on the subdomain interface. As a result, the mathematical theory of substructuring domain decomposition methods, which is grounded on the object decomposition of functions and trace theorems, cannot be readily applied. In fact, it is easy to find examples that prove that the standard BDDC method breaks down for unfitted meshes.

4.1. **Breakdown example.** Let us consider the problem in Fig. 4 and the BDDC(c) preconditioner in 2D. The left side of the domain is at distance ϵ to the interface. The underlying PDE is a Poisson problem with Neumann boundary conditions on the left (moving) side and Dirichlet Boundary conditions on the rest of the boundary, defined on a 2D rectangular domain that is dicretized with an unfitted mesh of bilinear Lagrange elements and partitioned into eight uniform subdomains. The example is designed in such a way that, as ε tends to zero, the area of subdomains ω^1 and ω^5 tend to zero (see Fig. 4a). Now, we consider a function $v \in \tilde{\mathbb{V}}$ that vanishes everywhere but in ω^1, where it has value 1 on all DOFs that belong to the edge shared by ω^1 and ω^2. It is clear that 1) the energy before weighting goes to zero with ε and 2) the energy after weighting is bounded below independently of ε after weighting. As a result,

$$\lambda_{\max} = \max_{\tilde{v} \in \tilde{\mathbb{V}}} \frac{\langle \mathcal{A}\mathcal{Q}\tilde{v}, \mathcal{Q}\tilde{v} \rangle}{\langle \mathcal{A}\tilde{v}, \tilde{v} \rangle} \tag{8}$$

tends to ∞ as $\varepsilon \to 0$. As a result, the condition number of the BDDC preconditioner cannot be uniform with respect to the surface intersection. (Similar examples can be designed for 3D problems and the other variants of the preconditioner.) It is clearly observed that the usual topological weighting operator leads to a solver that is not robust with respect to the position of the interface. It is also observed experimentally in Fig. 4 that the iteration count explodes as ε tends to zero.

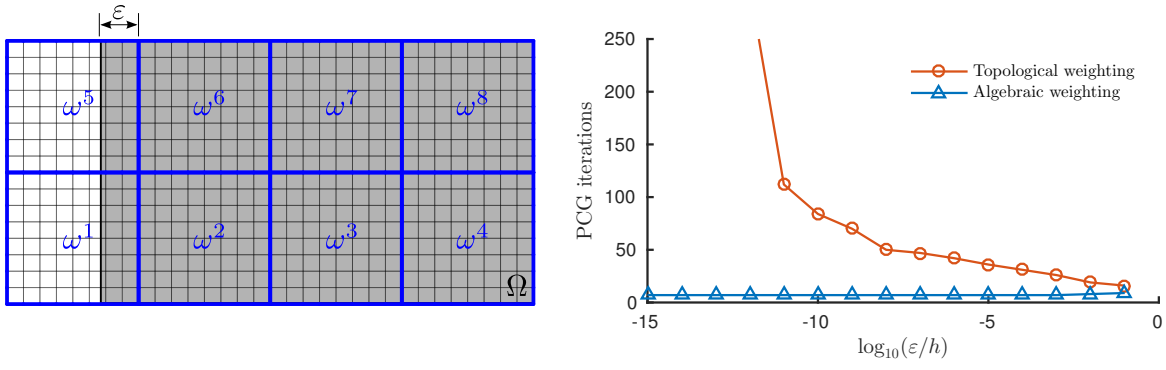

(a) (b)

Figure 4. Algebraic weighting vs topological weighting.

4.2. **Enlarged coarse spaces and adaptive BDDC methods.** In order to fix the BDDC preconditioner for unifitted FE meshes, one could consider as coarse corner DOFs all the nodes that violate (5) and are on the interface. Such approach would lead to a huge space. The authors note that reduced coarse spaces could be considered, especially when the cut cells are related to Neumann boundary conditions. In any case, even though improved solutions can be designed, the coarse space is still too large for practical computations. In the following section, we consider a different approach related to stifness weighting.

Recently, adaptive BDDC methods have been developed based on the idea to use specral solvers to define the coarse space in such a way that the bound for (8) can be controlled by the user (see, e.g., [39, 43]). The analysis of these methods does not rely on trace theorems (a difference with respect to classical BDDC methods). The analysis is algebraic (not relying on trace theorems, functional analysis tools, or FE inverse inequalities) and these algorithms could readily be used for unfitted FE methods to lead to provably robust solvers. However, such approach has two limitations: 1) the additional cost related to the eigenvalue solvers is very large and it only pays the price in some cases, and 2) it would lead to a severe load imbalance (only subdomains intersected by the surface would really require these eigenvalue solvers for many problems of interest). For these reasons, we do not pursue this line in this work.

4.3. **Stifness weighting operator.** In this section, we will present a modification of the BDDC preconditioner that has been experimentally observed to be robust with respect to small cuts in unfitted FE methods with a moderate increase of the coarse solver size. The use of stiffness matrix-based weighting has previously been used in other situations [20] with good results. Unfortunately, there are no rigorous condition number bounds for this approach.

We can represent functions in $\bar{\mathbb{V}}$ as vectors in $\mathbb{R}^{|\mathcal{N}|}$, and define the symmetric positive definite matrix $A_{ab} \doteq \langle \mathcal{A}\phi_b, \phi_a \rangle$ (analogously for the subassembled subdomain matrix A^{ω}_{ab}). We introduce the so-called *algebraic* weighting operator $\mathcal{W} : \mathbb{V} \to \bar{\mathbb{V}}$ for a given function $u \in \mathbb{V}$ and a node ξ of the mesh:

$$\mathcal{W}u(\xi) \doteq \frac{\sum_{\omega \in \text{neigh}(\xi)} A^{\omega}_{\xi\xi} u_{\omega}(\xi)}{\sum_{\omega \in \text{neigh}(\xi)} A^{\omega}_{\xi\xi}}, \tag{9}$$

Formula (9) provides the node values after the weighting, which uniquely defines a FE function in $\bar{\mathbb{V}}$. This type of weighting is referred to as the *algebraic* weighting since it is defined from the problem matrices, in contrast to the topological weighting (see Sect. 3.3) that is only based on the topology of the partition.

As we will show, the algebraic weighting leads to a much more robust BDDC method with respect to the relative position of domain boundary $\partial\Omega$ and the unfitted grid. We consider the 2D example in Fig. 4, for which the BDDC preconditioner with topological weighting is not robust. Fig. 4b shows the number of PGC iterations in function of the position of the interface. It is clearly observed that the usual topological weighting operator leads to a solver that is not robust with respect to the position of the interface. The iteration count explodes as ε tends to zero. On the other hand, the algebraic weighting operator leads to an extremely robust solver. The number of iterations are nearly independent of the position of the interface, even when the value ε/h reaches machine precision.

Using the algebraic weighting, we have successfully solved also complex 3D examples (see the numerical results in Sect. 5). However, we have observed that the number of iterations needed to solve the systems for meshes with cut elements are (about 2 times) higher that the iterations needed to solve the same problem without cutting the mesh elements. This increase in iterations is not dramatic, but further improvement is possible. To this end, we propose two further enhancements of the BDDC method based on the definition of the coarse objects that lead to almost identical iteration counts than the case in which all cut cells are filled and the standard BDDC method is used.

4.4. **New coarse objects: variant 1.** The current mathematical analysis of BDDC methods is for body-fitted meshes, and therefore, it implicitly assumes that the coarse objects are not intersected by the boundary $\partial\Omega$. However, when body-fitted meshes are allowed, the objects can be cut by $\partial\Omega$ (see Fig. 5a). In this situation, the main mathematical properties required to proof scalability of the

BDDC method are lost. It can lead to situations in which the corner constraints are exterior corners of cells with close to zero η_e, possibly leading to close to singular local problems. Motivated by this fact, we propose a modification in the definition of the coarse objects, which eliminates intersections with the boundary $\partial\Omega$ even if unfitted grids are considered. Since fixing the rigid-body modes for all the edges is enough in 3D (see, e.g., [31]), we will only consider the modification of edge coarse DOFs.

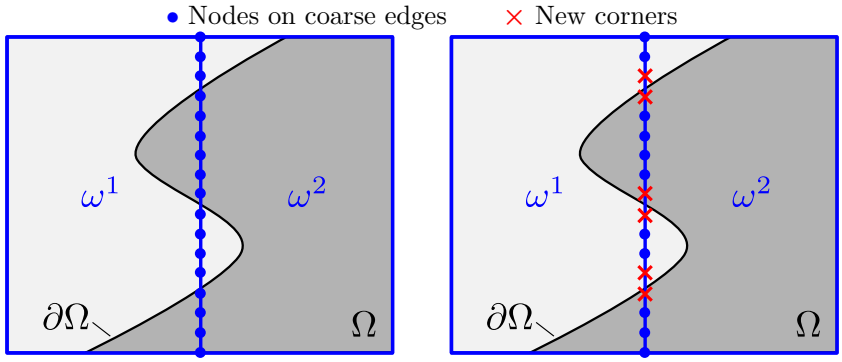

(a) Before (b) After

FIGURE 5. Splitting a cut edge into several uncut edges (first variant).

The new objects are computed in two main steps. The first one is to build the standard objects Λ_O of the BDDC method for θ_{act}, as presented in Sect. 3.2. The second step is to identify the edges in Λ_O cut by the boundary and split them into new edges/corners, leading to uncut edges (see Fig. 5). We consider a segregation of a coarse edge $\lambda \in \Lambda_E$ as follows. Given the set of FE edges $t \subset \lambda$, we segregate them into interior edges $t \subset \Omega$, exterior edges such that $t \cap \Omega = \emptyset$, and the remaining cut edges. All DOFs on cut edges are considered as coarse corner DOFs. The rest of FE edges are agregated into *connected* subedges (see Fig. 5b). This procedure is done efficiently by using the data structures commonly available in domain decomposition codes and immersed boundary methods. It requires to perform a loop on the edges objects (which are usually identified anyway in a domain decomposition code), and determine if those edges are intersected by the interface or not. Intersection points are usually already pre-computed in immersed boundary codes. As a result, no further significant overhead is introduced. Moreover, this operation has to be performed only in the subdomains cut by the boundary, which is typically a small set. The numerical examples (see Sect.5) show that by splitting edge objects in this way, optimal weak scaling is recovered if Neumann boundary conditions are imposed on the unfitted parts of the domain boundary $\partial\Omega$.

4.5. **New coarse objects: variant 2.** However, the case of weak imposition of Dirichlet conditions has been observed to be more challenging and requires further elaboration. In that case, we propose a second strategy to split the edges. The method is motivated by the fact that the existing domain decomposition theory (see Eq. (6)) assumes constant weighting coefficients within each coarse object in order to derive condition number bounds. We note that the coefficient β required to ensure coercivity in Nitsche's method might be very different between neighbor elements (even different orders of magnitude) in function of the location of the interface. This induces strong variations in the coefficients of the stiffness matrix, and therefore highly oscillatory weighting coefficients, when they are computed as in (9). As a result, the weighting coefficient might be non-constant within the BDDC objects cut by the interface (see Fig. 6a), which violates one of the basic assumptions of the domain decomposition analysis. Motivated by this fact, we try to build objects that have constant weighting coefficients. To this end, we follow two steps. First, we build the standard BDDC objects (Sect. 3.2). Then, we split all coarse objects classified as *edges* that have non-constant weighting coefficients. Precisely, we aggregate neighbor nodes on the edge that have the same weighting coefficient, and then we set the remaining un-aggregated nodes as coarse corners (see Fig. 6b). Since edges are one dimensional curves,

this aggregation can be easily performed. Moreover, this process is performed only for the small set of subdomains that are cut by the boundary. By modifying the coarse objects in this way, an optimal weak scaling is also recovered for weak Dirichlet boundary conditions, as showed in the following 3D numerical examples.

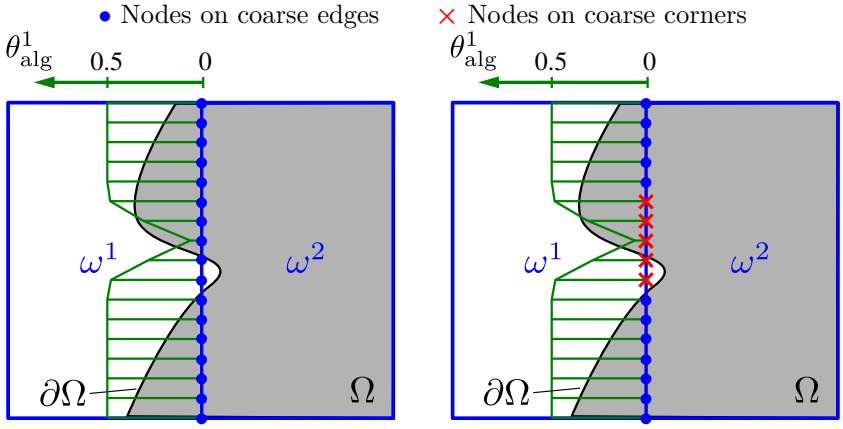

(a) Before (b) After

FIGURE 6. Splitting a cut edge into several uncut edges (second variant).

5. NUMERICAL EXPERIMENTS

5.1. **Setup.** In the numerical examples below, we consider the Poisson problem with the discretization being used in Sect. 2.2 defined on five different complex 3D domains shown in Fig. 7. The chosen geometries are 1) a sphere, 2) a body that reminds to a popcorn flake, 3) a hollow block, 4) a 3-by-3-by-3 array of such blocks, and 5) a spiral. These geometries are used often in the literature to study the performance of unfitted FE methods (see, e.g., [15]). This variety of shapes in considered here to evaluate the versatility of the proposed solvers and their capacity to deal with complex and diverse 3D domains. It is basic, since the solvers proposed are based on heuristic motivations, and numerical experimentation is the only way to prove their robustness.

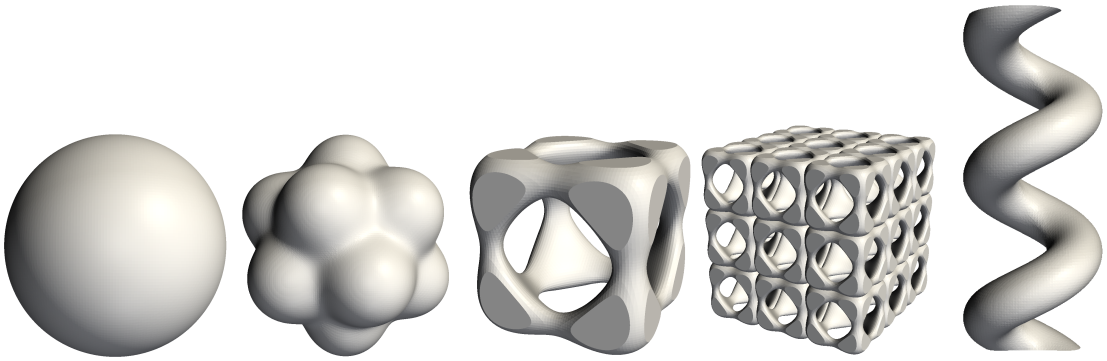

(a) Sphere. (b) Popcorn flake. (c) Block. (d) Array of blocks. (e) Spiral.

FIGURE 7. Geometries used in the numerical examples.

For each geometry, weak scalability tests are executed in order to evaluate the performance of the solvers. To this end, a family of computational meshes are generated for each piece spanning from coarse meshes up to fine meshes with several millions of elements (see Table 1 for details). Note that

the subdomain partitions are refined at the same rate than the meshes, leading to a constant ratio H/h, which should lead to an optimal weak scaling of the BDDC preconditioner. That is, the number of iterations of the underlying preconditioned conjugate gradient solver is expected to be (asymptotically) independent of the problem size. The main purpose of the numerical example below is to check that the proposed BDDC preconditioners for non-conforming grids have this optimal performance. In the remainder of this sub-section, we further detail the setup of the numerical examples.

TABLE 1. Information of the computational meshes used in the examples.

	Id.	n^{sd}	n^{el}	n^{sdel}	n^{dof}	n^{sddof}	h	H	H/h
Sphere	1	8	1064	512	1461	729	0.1250	1.0000	8.0
	2	32	7160	512	8577	729	0.0625	0.5000	8.0
	3	160	51968	512	57123	729	0.0312	0.2500	8.0
	4	1064	396040	512	415775	729	0.0156	0.1250	8.0
	5	7160	3090104	512	3167421	729	0.0078	0.0625	8.0
Popcorn flake	1	8	1920	512	2559	729	0.1125	0.9000	8.0
	2	60	12936	512	15181	729	0.0563	0.4500	8.0
	3	324	95256	512	103701	729	0.0281	0.2250	8.0
	4	1924	729768	512	762497	729	0.0141	0.1125	8.0
	5	12936	5708900	512	5837623	729	0.0070	0.0563	8.0
Block	1	8	2120	512	3280	729	0.2312	1.8500	8.0
	2	56	14336	512	18772	729	0.1156	0.9250	8.0
	3	304	103152	512	120164	729	0.0578	0.4625	8.0
	4	2120	776584	512	842652	729	0.0289	0.2312	8.0
	5	14336	6011832	512	6272228	729	0.0145	0.1156	8.0
Array	1	216	53904	512	72796	729	0.2271	1.8167	8.0
	2	1512	376696	512	451516	729	0.1135	0.9083	8.0
	3	8208	2747880	512	3064716	729	0.0568	0.4542	8.0
Spiral	1	2	435	512	838	729	0.1500	1.2000	8.0
	2	16	2277	512	3501	729	0.0750	0.6000	8.0
	3	96	14309	512	18511	729	0.0375	0.3000	8.0
	4	441	99991	512	115378	729	0.0188	0.1500	8.0
	5	2277	744170	512	802975	729	0.0094	0.0750	8.0
	6	14309	5733813	512	5963098	729	0.0047	0.0375	8.0

Legend
n^{sd}: number of subdomains
n^{el}: total number of elements
n^{sdel}: max. number of elements in a subdomain
n^{dof}: total number of DOFs
n^{sddof}: max. number of DOFs in a subdomain
h: element size, H: subdomain size.

We define the source term and boundary conditions of the Poisson equation such that the PDE has an exact solution, namely

$$u(x,y,z) = \sin\left(5\pi\left(x^2+y^2+z^2\right)^{1/2}\right), \quad (x,y,z) \in \Omega \subset \mathbb{R}^3.$$

We use the exact solution to assess the quality of the numerical approximations provided by the solvers. In order to define the boundary conditions, we partition the boundary $\partial\Omega$ into two non-overlapping sets Γ_{bb} and Γ_{cut} (see Fig. 8). The set Γ_{bb} is the intersection of $\partial\Omega$ with the bounding box used to define the computational mesh, and Γ_{cut} is its complement, $\Gamma_{\text{cut}} \doteq \partial\Omega - \Gamma_{\text{bb}}$. After discretizing the bounding box with a Cartesian mesh, the set Γ_{bb} results to be aligned with the mesh faces. On the contrary, Γ_{cut} is a non-conforming boundary living inside the cut elements. We define strong Dirichlet conditions on the boundary Γ_{bb}. Since Γ_{bb} is aligned with the mesh, strong Dirichlet boundary conditions are imposed by constraining the corresponding mesh nodes. On the other hand, we impose two different types of boundary conditions on the unfitted surface Γ_{cut}, either 1) weak Dirichlet boundary conditions

using Nitsche's method or 2) Neumann boundary conditions. Note that for the popcorn flake and the spherical domain, the set Γ_{bb} is empty since all the boundary $\partial\Omega$ in non-conforming (see Fig. 8). In these cases, we strongly impose the value of the solution at an interior node of Ω in order to have a well posed problem when considering Neumann conditions on the entire boundary.

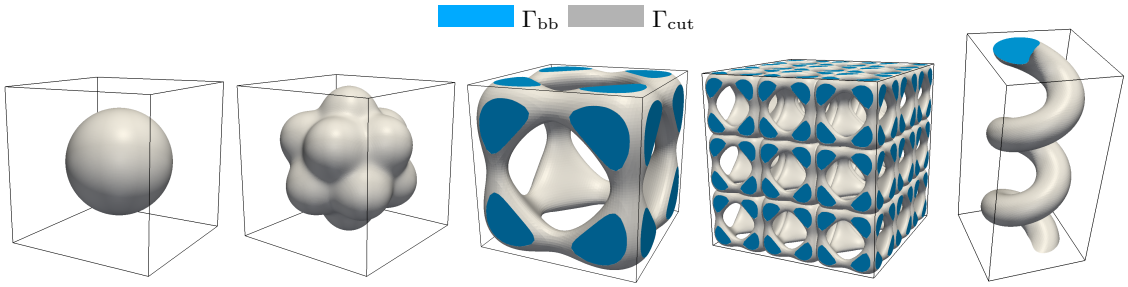

(a) Sphere. (b) Popcorn flake. (c) Block. (d) Array of blocks. (e) Spiral.

FIGURE 8. Illustration of the sets Γ_{bg} and Γ_{cut}.

After the discretization process, the resulting systems of linear equations $Ax = b$ are solved with a BDDC-preconditioned conjugate gradient (PCG) solver. Convergence of the PCG solver is declared when the 2-norm of the algebraic residual $r \doteq b - Ax$ is smaller that the 2-norm of the right-hand-side b times a tolerance, namely $|r|_2 < \text{tol} \cdot |b|_2$. In all cases, the selected tolerance is $\text{tol} = 10^{-9}$. With this value, the error committed in solving the linear system is negligible in front of the discretization error associated with the FE mesh. This is demonstrated by the fact that the optimal convergence rate of FEM is obtained by using this tolerance for all the geometries (see Fig. 9).

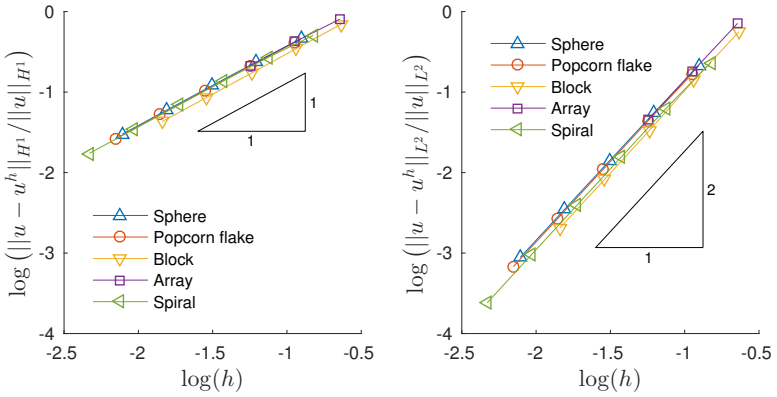

FIGURE 9. Convergence of the discretization error in H^1 and L^2 norms.

Six different versions of the BDDC preconditioner are studied in the examples, see Table 2. They differ in 1) the definition of the coarse objects (either the Standard definition of section 3.2 or the extended definitions for non-conforming meshes in Sects. 4.4 and 4.5) and 2) the subset of objects used to define the coarse DOFs (either corners and edges, or corners edges and faces). In all the cases, the algebraic weighting operator of Sect. 4.3 is considered. We do not include the topological weighting operator (detailed in Sect. 3.3) since it leads to a solver that is not robust with respect to the relative position between the geometry and the mesh (see, e.g., the breakdown example in Sect. 4.1). We have experimentally observed that the topological weighting operator leads to unacceptably large iteration counts for the 3D geometries studied here. In particular, more than 1000 iterations are needed to solve the "popcorn flake" with a medium-size mesh (mesh # 3 described in Table 1) for the case of Neumann boundary conditions on Γ_{cut}.

Finally, in order to compare the performance of the proposed BDDC methods for unfitted grids, with respect to the optimal behavior of BDDC for conventional conforming meshes, we construct auxiliary meshes by replacing the cut elements by standard full elements (see Fig. 10). Our goal is to recover a similar performance for cut elements, that the one obtained with full elements.

TABLE 2. Description of the BDDC preconditioners used in the numerical examples.

Identifier	Object definition	Selected objects	Weighting operator
Standard BDDC(cef)	As in Sect. 3.2	corners, edges, faces	As in Sect. 4.3
Standard BDDC(ce)	As in Sect. 3.2	corners, edges	As in Sect. 4.3
1st variant BDDC(cef)	As in Sect. 4.4	corners, edges, faces	As in Sect. 4.3
1st variant BDDC(ce)	As in Sect. 4.4	corners, edges	As in Sect. 4.3
2nd variant BDDC(cef)	As in Sect. 4.5	corners, edges, faces	As in Sect. 4.3
2nd variant BDDC(ce)	As in Sect. 4.5	corners, edges	As in Sect. 4.3

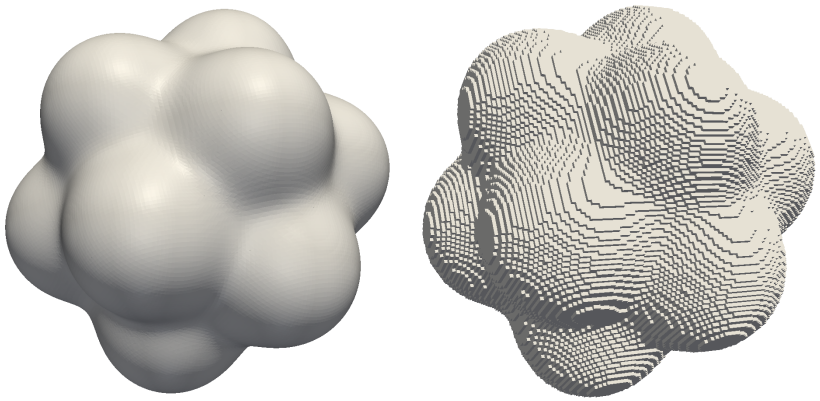

(a) Cut elements. (b) Full elements.

FIGURE 10. The "popcorn flake" geometry reconstructed with cut elements and with full elements.

5.2. **"Popcorn flake" example.** For the sake of clarity, we first present the numerical results for a single geometry, namely the "popcorn flake" geometry in Fig. 7b. The results for the other shapes are discussed in Sect. 5.3 below. Fig. 11 shows the result of the weak scaling test for this geometry, where Fig. 11a is for the case of Neumann boundary conditions on Γ_{cut}, and Fig. 11b is for weak Dirichlet boundary conditions. It is observed that the standard BDDC method leads to a perfect weak scaling for the case of full elements (as expected) for the two types of boundary conditions. The performance of the standard BDDC method when dealing with cut elements (i.e., red lines in Fig. 11) is decent, in the sense that the number of iterations increases only mildly with the problem size, but it is clear that the absolute number of iterations needed to solve the problems is more than twice the iterations required for full elements (which is the target value). In contrast, both the 1st and 2nd variants of BDDC achieve a perfect weak scaling with iteration counts close to the ones associated with full elements for the case of Neumann boundary contitions (see blue and yellow lines in Fig. 11a).

However, for Dirichlet boundary conditions, only the 2nd variant provides a perfect weak scaling close to the one observed for full elements (see yellow line in Fig. 11b). In conclusion, the 1st variant of the proposed method is enough to achieve a perfect weak scaling for Neumann boundary conditions, whereas the 2nd variant is required for Dirichlet conditions imposed with Nitsche's method. Similar

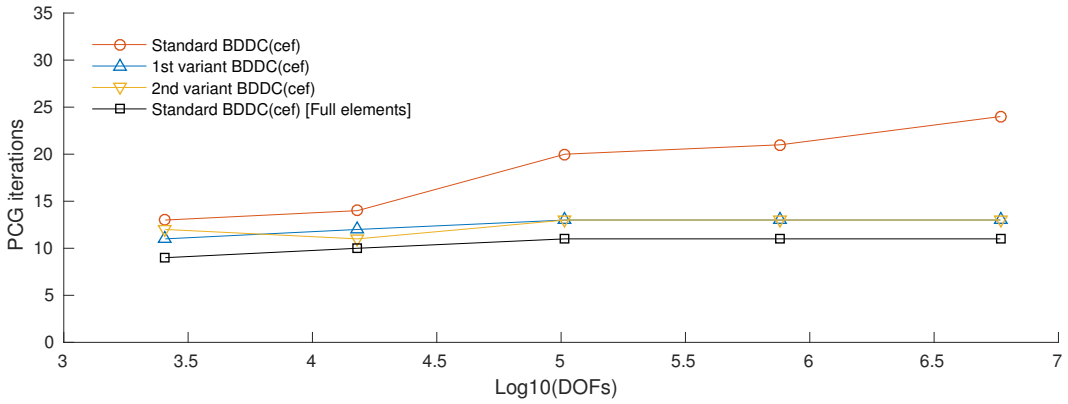

(a) Results for Neumann boundary conditions.

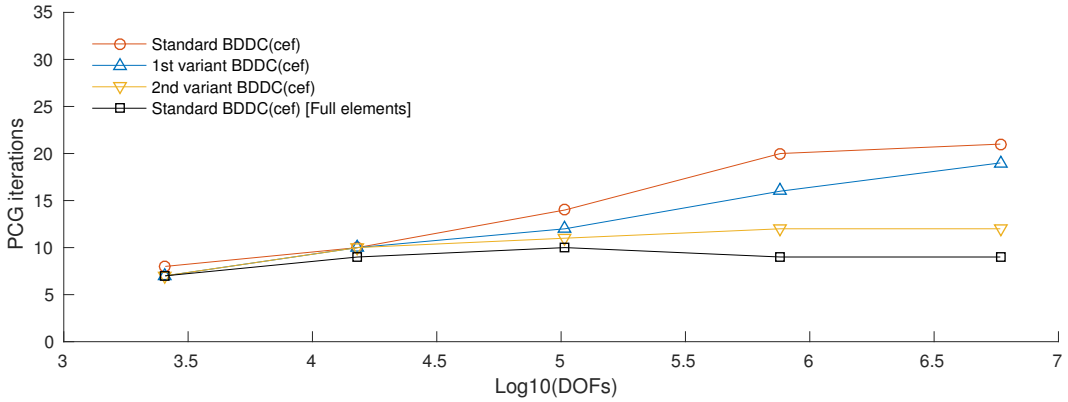

(b) Results for Dirichlet boundary conditions.

FIGURE 11. "Popcorn flake" example: Weak scalability test.

results are obtained if the coarse space is defined with corners, edges, and faces (cef), or only with corners and edges (ce) (see Fig. 12). The main difference is that the iterations associated with corners and edges are slightly larger than for corners, edges, and faces, which is the expected behavior in BDDC methods. The iteration count is smaller in the later case since the associated coarse space is larger.

Fig. 13 shows the size of the coarse space of the new BDDC methods, namely variants 1 and 2, with respect to the size of the standard coarse space. Fortunately, the size of the coarse spaces associated with variants 1 and 2 tends to be equal to the size of the standard coarse space in BDDC as the mesh is refined (see Fig. 13). This behavior is explained by the fact that extra corners are only added on edges that are close to cut elements. Since the number of BDDC objects in the interior of Ω scales much faster that the objects cut by the boundary $\partial\Omega$, the number of extra corners added near the boundary become nearly negligible as the partition is refined. As a result, the finer is the mesh the more competitive are the proposed BDDC preconditioners for unfitted meshes. Thanks to this behavior and for the perfect weak scaling results in [5, 7], the proposed BDDC methods (variants 1 and 2) are suitable for large-scale problems requiring fine meshes and a large number of subdomains. For Neumann boundary conditions, variant 1 is the best choice since it leads to a smaller coarse space than variant 2 for approximately the same number of PCG iterations. For Dirichlet boundary conditions with Nitsche's method, the 2nd variant is the preferred choice specially for fine meshes, since it leads to a better weak scaling.

It is also desirable that the corner detection mechanisms in variants 1 and 2 do not have a negative impact in the load balancing of the local operations involved in the BDDC preconditioner, namely

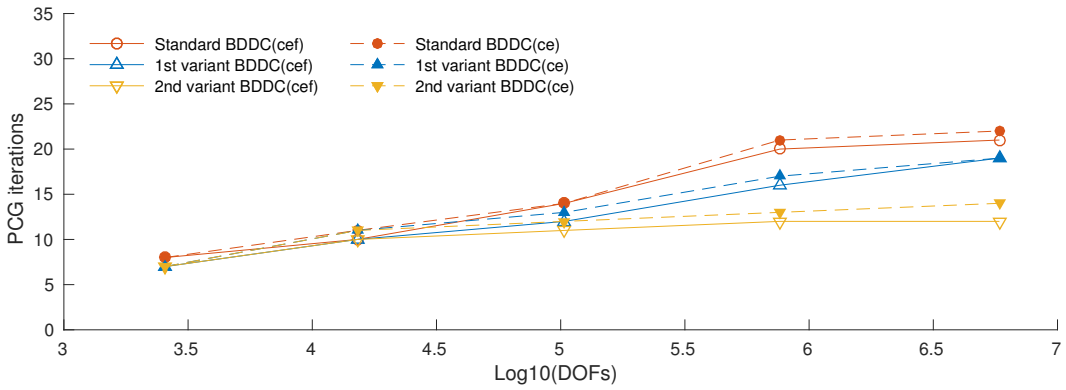

FIGURE 12. "Popcorn flake" example: Weak scalability test. Comparison between BDDC(cef) and BDDC(ce).

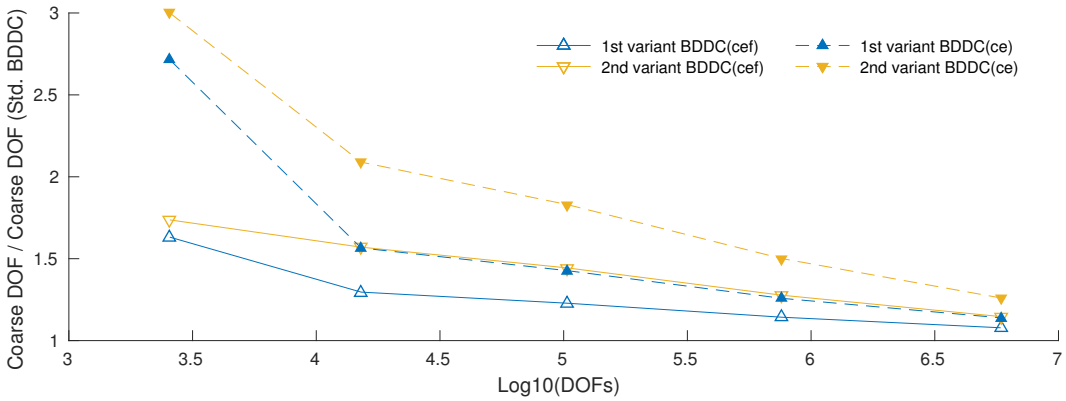

FIGURE 13. "Popcorn flake" example: Size of the coarse space for the new BDDC methods (variants 1 and 2 in Table 2) with respect the size of the usual coarse space in BDDC.

the number of local Neumann problems required to compute the coarse basis functions, which is the main algorithmic source of load imbalance (the number of local Dirichlet problems and constrained Neumann problems in the application step are identical for all subdomains). The effect of the proposed method on the load balancing is studied in Fig. 14. Here, the extra extra corners added by 1st and 2nd variants of the new method can potentially affect the load balancing. However, the balancing also becomes better as the mesh is refined (see Fig 14). That is, the potential negative effect of the extra corners on the load balancing is small for fine meshes. The results in Fig. 14 are for the 2nd variant of the BDDC(ce) method. Even better results are obtained for the 1st variant and the BDDC(cef) method. The plots are not included for the sake of brevity since they do not add extra information.

In a final phase of the experiment, we study the robustness of the methods with respect to the position of the cut geometry. To this end, we slightly move the location of the background mesh by an arbitrary value ε in one of the three spatial directions. By translating the mesh in this way, the location of the geometry changes and creates different kinds of intersections with the mesh elements, leading to cut elements with different active volume fractions that can affect the conditioning of the problem in different ways. Our goal is to show that the proposed methods are not affected by this translation and the different configurations of cut elements. The values for ε are chosen in the closed interval $[-h, +h]$, where h is the element size of the background mesh. Fig. 15 shows the number of

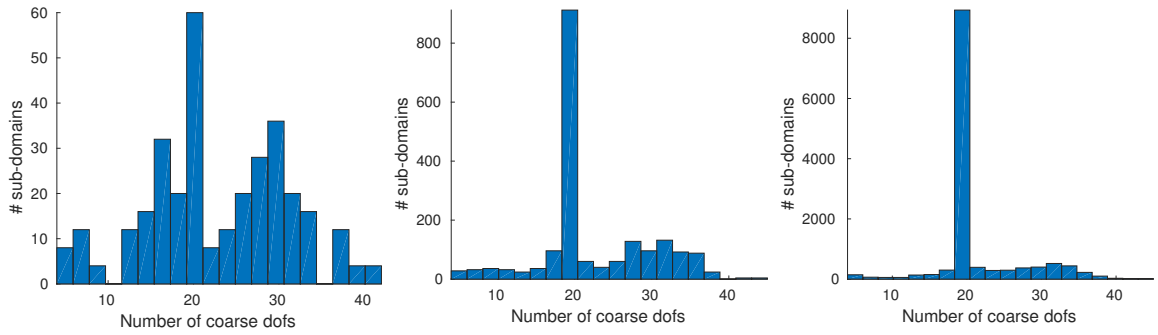

(a) Mesh 3. (b) Mesh 4. (c) Mesh 5.

FIGURE 14. "Popcorn flake" example: Distribution of the number of coarse basis functions across the subdomains.

iterations of the PCG solver in function of the value ε, for the mesh number 3 in Table 1. The results are for weak Dirichlet boundary conditions, which is the most challenging case. The results in Fig. 15 show that the 2nd variant is very robust with respect to the position of the interface. The maximum change in the number of iterations is only 1 iteration for all the 51 different values of ε considered. The other methods have also a decent robustness with a maximum change of 4 iterations.

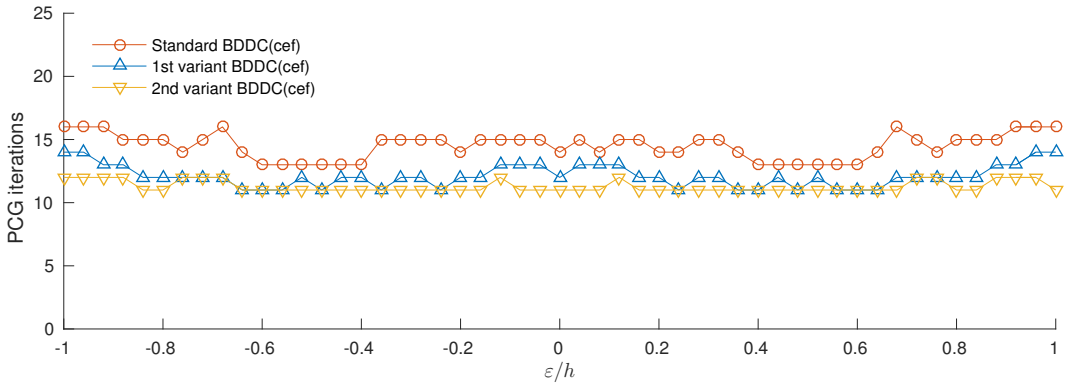

FIGURE 15. "Popcorn flake" example: Study of the influence of the position of the boundary on the porposed methods.

5.3. **Multi-body example: sphere, popcorn flake, block, array of blocks and spiral.** The purpose of this final example is to show that the proposed BDDC preconditioners are able to deal with different complex geometries, and that the results obtained in previous example are qualitatively reproduced also in other settings. To this end, we study the BDDC preconditioners for all the geometries previously presented in Fig. 7. The results of the weak scalability test are presented in Fig. 16. The figure shows that performance of the methods for this set of geometries is similar to the one previously obtained for the "popcorn flake". For Neumann boundary conditions a perfect weak scaling is obtained for both the 1st and 2nd variants, whereas for Dirichlet boundary conditions a perfect weak scaling is only obtained with the 2nd variant. It is remarkable that the number of iterations is almost identical for all the geometries for the finest meshes (see 1st and 2nd variants for Neumann conditions, and 2nd variant for Dirichlet conditions). This is true even for the array of blocks, which is the most complex geometry considered herein.

The size of the coarse spaces behaves similar to the previous example (see Fig. 17). The increment in the coarse space size associated with variants 1st and 2nd becomes smaller in front of the size of the

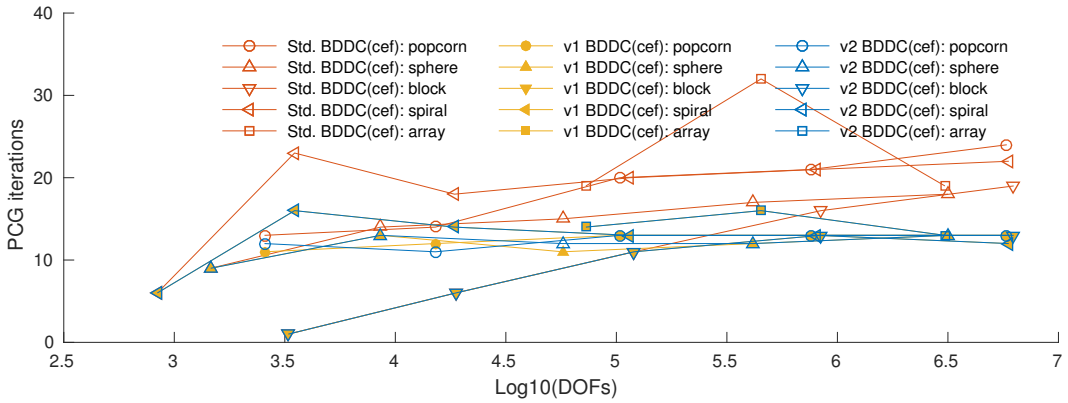

(a) Results for Neumann boundary conditions.

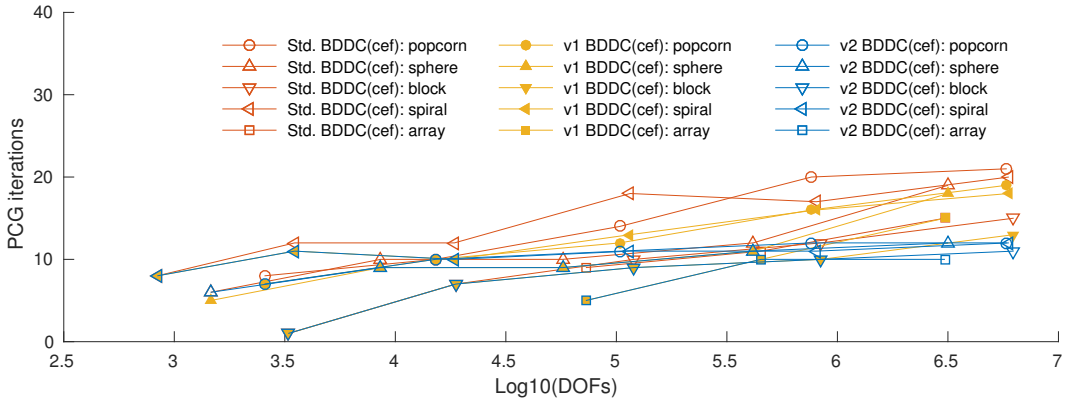

(b) Results for Dirichlet boundary conditions.

FIGURE 16. Multi-body example: Weak scalability test for all the different geometries displayed in Fig. 7.

coarse space in standard BDDC as the mesh is refined. This is again due to the fact that the number of objects in the bulk of the domain scales faster than the objects cut by the boundary. For that reason, the sphere is the shape that presents a smaller increment in the coarse space size. It minimizes the cut surface per unit of domain volume. Note, however, that the coarse space size of the variants 1 and 2 tends also to the coarse space size of standard BDDC for the other shapes (see, e.g., the spiral or the array of blocks).

Finally, we study the robustness of the methods with respect to the position of the unfitted boundary $\partial\Omega$. In all cases, we move the background mesh a value $\varepsilon \in [-h, +h]$ to see the influence of the boundary location. Fig. 18 shows the number of iterations of the PCG solver with respect to the value ε for all the shapes. The results correspond to weak Dirichlet boundary conditions, which is the more challenging case. The results are similar to previous example. It is observed that the performance of the 2nd variant is very robust. It is almost independent of the position of the interface with a maximum change of 1 iteration in all the considered geometries.

In summary, the presented results show that the proposed variants of the BDDC method are able to provide excellent results for different and complex geometries discretized with non-conforming FE meshes. As in previous example, the 1st variant is already enough to achieve a perfect weak scaling for all the shapes for the case of Neumann boundary conditions. The 2nd variant is the only one that leads to perfect weak scaling for Dirichlet conditions. It is remarkable that the number of PCG iterations is nearly independent from the geometry for the finest meshes. This shows that the methods are able to effectively take rid of the complexities introduced by the complex geometries and cut elements. The

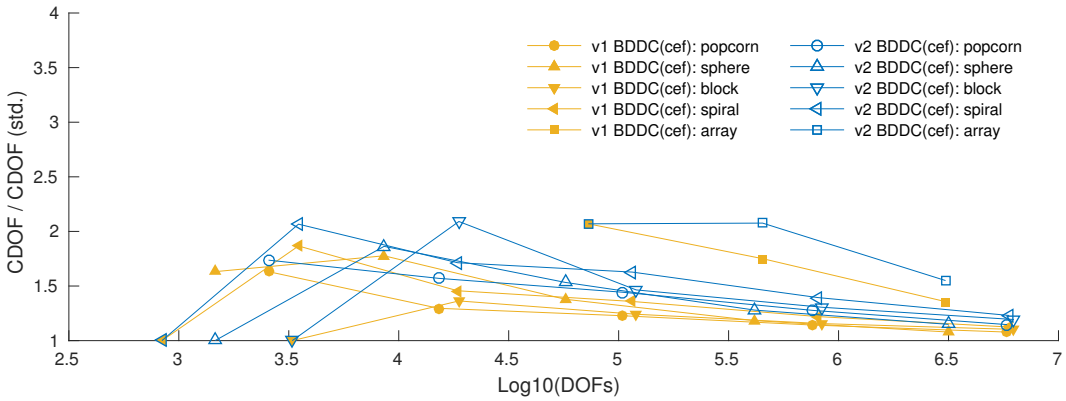

FIGURE 17. Multi-body example: Size of the coarse space for the new BDDC methods (variants 1 and 2 in Table 2) with respect to the size of the usual coarse space in BDDC.

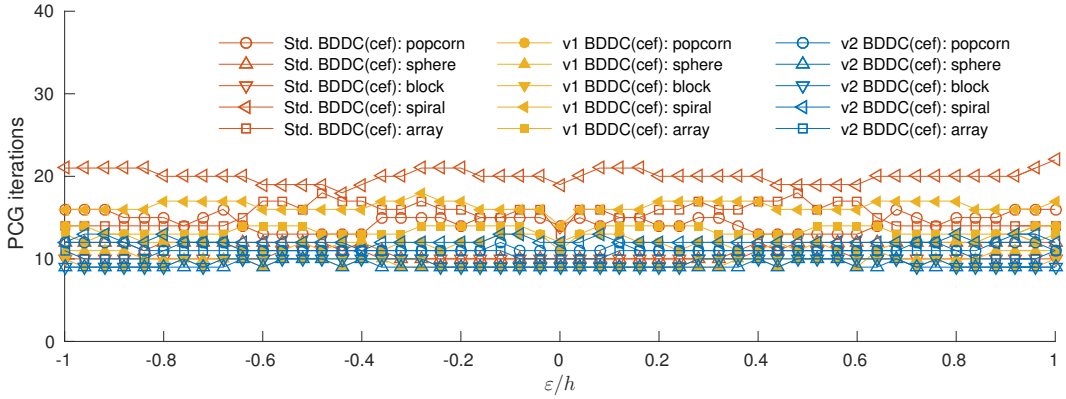

FIGURE 18. Multi-body example: Study of the influence of the position of the boundary on the porposed methods.

size of the coarse spaces of the new methods tends to the size for standard BDDC as the mesh is refined. Therefore, the efficiency of the new methods is comparable to the one of standard BDDC for fine meshes.

6. CONCLUSIONS

In this work, we have analyzed the main problems related to the use of BDDC methods to unfitted FE meshes. From a mathematical point of view, there are different assumptions in the analysis of these algorithms that do not hold anymore for unfitted meshes. As a result, the polylogarithmic condition number bounds in BDDC methods are not true in general. In fact, using a breakdown example, we have observed that the condition number of these methods cannot be robust with respect to the intersection between the surface and the background mesh. One can fulfill again the assumptions of the BDDC analysis by considering as corner constraint all the DOFs on the interface that belong to cut cells. Unfortunately, even thought the resulting algorithm is perfectly robust, such approach would lead to a dramatic increase in the coarse problem size, and it is not considered in practice.

Alternatively, and motivated by both numerical experimentation and the mathematical analysis of BDDC methods, we have considered 1) a stiffness-based weighting and 2) a subdivision of the BDDC edges. The first ingredient has turned to be an essential ingredient to get a robust BDDC solver for unfitted meshes. On the other hand, the enrichment of the coarse BDDC space is especially

important in regions with weak imposition of Dirichlet boundary conditions. As usual, the use of stiffness weighting is not backed up with a full mathematical theory. Instead, a complete set of numerical experiments has been performed to show the robustness and algorithmic weak scalability of the proposed solvers.

Future work will include the extension to more complex (non-coercive) physical problems and to ghost penalty stabilization. The use of stabilized versions is expected to be easier than the problem considered herein, since the condition number bound in these cases can be bounded independently of the location of the intersection. Furthermore, we want to implement these solvers within the `FEMPAR` library [1, 6], in order to show weak scalability results for its highly scalable implementation [4, 5] of multilevel BDDC preconditioners. We also want to apply these methods to adaptive octree meshes and space-filling curves [2].

References

[1] *FEMPAR: Finite element multiphysics parallel solvers*. `https://gitlab.com/fempar/fempar`.

[2] M. Bader, *Space-Filling Curves: An Introduction With Applications in Scientific Computing*, Springer Science & Business Media, Oct. 2012. Google-Books-ID: eIe_OdFP0WkC.

[3] S. Badia, *On stabilized finite element methods based on the Scott–Zhang projector. Circumventing the inf–sup condition for the Stokes problem*, Computer Methods in Applied Mechanics and Engineering, 247–248 (2012), pp. 65–72.

[4] S. Badia, A. F. Martín, and J. Principe, *Implementation and Scalability Analysis of Balancing Domain Decomposition Methods*, Archives of Computational Methods in Engineering, 20 (2013), pp. 239–262.

[5] S. Badia, A. F. Martín, and J. Principe, *Multilevel Balancing Domain Decomposition at Extreme Scales*, SIAM Journal on Scientific Computing, (2016), pp. C22–C52.

[6] S. Badia, A. F. Martín, and J. Principe, *FEMPAR: An object-oriented parallel finite element framework*, In preparation, (2017).

[7] S. Badia, A. F. Martín, and J. Principe, *A Highly Scalable Parallel Implementation of Balancing Domain Decomposition by Constraints*, SIAM Journal on Scientific Computing, 36 (2014), pp. C190–C218.

[8] S. Badia, F. Nobile, and C. Vergara, *Fluid–structure partitioned procedures based on Robin transmission conditions*, Journal of Computational Physics, 227 (2008), pp. 7027–7051.

[9] W. Bangerth, C. Burstedde, T. Heister, and M. Kronbichler, *Algorithms and data structures for massively parallel generic adaptive finite element codes*, ACM Transactions on Mathematical Software (TOMS), 38 (2011), p. 14.

[10] R. Becker, *Mesh adaptation for Dirichlet flow control via Nitsche's method*, Communications in Numerical Methods in Engineering, 18 (2002), pp. 669–680.

[11] T. Belytschko, N. Moës, S. Usui, and C. Parimi, *Arbitrary discontinuities in finite elements*, International Journal for Numerical Methods in Engineering, 50 (2001), pp. 993–1013.

[12] L. Berger-Vergiat, H. Waisman, B. Hiriyur, R. Tuminaro, and D. Keyes, *Inexact Schwarz-algebraic multigrid preconditioners for crack problems modeled by extended finite element methods*, International Journal for Numerical Methods in Engineering, 90 (2012), pp. 311–328.

[13] S. C. Brenner and R. Scott, *The Mathematical Theory of Finite Element Methods*, Springer, softcover reprint of hardcover 3rd ed. 2008 ed., Nov. 2010.

[14] E. Burman, *Ghost penalty*, Comptes Rendus Mathematique, 348 (2010), pp. 1217–1220.

[15] E. Burman, S. Claus, P. Hansbo, M. G. Larson, and A. Massing, *CutFEM: Discretizing geometry and partial differential equations*, International Journal for Numerical Methods in Engineering, 104 (2015), pp. 472–501.

[16] E. Burman and P. Hansbo, *Fictitious domain finite element methods using cut elements: II. A stabilized Nitsche method*, Applied Numerical Mathematics, 62 (2012), pp. 328–341.

[17] C. Burstedde and J. Holke, *A Tetrahedral Space-Filling Curve for Nonconforming Adaptive Meshes*, SIAM Journal on Scientific Computing, 38 (2016), pp. C471–C503.

[18] C. Burstedde, L. C. Wilcox, and O. Ghattas, *p4est : Scalable Algorithms for Parallel Adaptive Mesh Refinement on Forests of Octrees*, SIAM Journal on Scientific Computing, 33 (2011), pp. 1103–1133.

[19] F. de Prenter, C. V. Verhoosel, G. J. van Zwieten, and E. H. van Brummelen, *Condition number analysis and preconditioning of the finite cell method*, Computer Methods in Applied Mechanics and Engineering, 316 (2017), pp. 297–327.

[20] C. R. Dohrmann, *A Preconditioner for Substructuring Based on Constrained Energy Minimization*, SIAM Journal on Scientific Computing, 25 (2003), pp. 246–258.

[21] M. Dryja, J. Galvis, and M. Sarkis, *BDDC methods for discontinuous Galerkin discretization of elliptic problems*, Journal of Complexity, 23 (2007), pp. 715–739.

[22] B. Gmeiner, M. Huber, L. John, U. Rüde, and B. Wohlmuth, *A quantitative performance study for Stokes solvers at the extreme scale*, Journal of Computational Science, 17, Part 3 (2016), pp. 509–521.

[23] J. Guzmán and M. Olshanskii, *Inf-sup stability of geometrically unfitted Stokes finite elements*, arXiv:1605.09681 [math], (2016). arXiv: 1605.09681.

[24] A. Hansbo and P. Hansbo, *An unfitted finite element method, based on Nitsche's method, for elliptic interface problems*, Computer methods in applied mechanics and engineering, 191 (2002), pp. 5537–5552.

[25] B. Hiriyur, R. Tuminaro, H. Waisman, E. Boman, and D. Keyes, *A Quasi-algebraic Multigrid Approach to Fracture Problems Based on Extended Finite Elements*, SIAM Journal on Scientific Computing, 34 (2012), pp. A603–A626.

[26] T. Isaac, C. Burstedde, L. C. Wilcox, and O. Ghattas, *Recursive Algorithms for Distributed Forests of Octrees*, arXiv:1406.0089 [cs], (2014). arXiv: 1406.0089.

[27] D. Kamensky, M.-C. Hsu, D. Schillinger, J. A. Evans, A. Aggarwal, Y. Bazilevs, M. S. Sacks, and T. J. R. Hughes, *An immersogeometric variational framework for fluid–structure interaction: Application to bioprosthetic heart valves*, Computer Methods in Applied Mechanics and Engineering, 284 (2015), pp. 1005–1053.

[28] G. Karypis, *METIS and ParMETIS*, in Encyclopedia of Parallel Computing, D. Padua, ed., Springer US, 2011, pp. 1117–1124.

[29] G. Karypis and V. Kumar, *A Fast and High Quality Multilevel Scheme for Partitioning Irregular Graphs*, SIAM Journal on Scientific Computing, 20 (1998), pp. 359–392.

[30] A. Klawonn, M. Lanser, and O. Rheinbach, *Toward Extremely Scalable Nonlinear Domain Decomposition Methods for Elliptic Partial Differential Equations*, SIAM Journal on Scientific Computing, 37 (2015), pp. C667–C696.

[31] A. Klawonn, O. B. Widlund, and M. Dryja, *Dual-Primal FETI Methods for Three-Dimensional Elliptic Problems with Heterogeneous Coefficients*, SIAM Journal on Numerical Analysis, 40 (2002), pp. 159–179.

[32] W. E. Lorensen and H. E. Cline, *Marching cubes: A high resolution 3D surface construction algorithm*, ACM SIGGRAPH Computer Graphics, 21 (1987), pp. 163–169.

[33] J. Mandel and C. R. Dohrmann, *Convergence of a balancing domain decomposition by constraints and energy minimization*, Numerical Linear Algebra with Applications, 10 (2003), pp. 639–659.

[34] O. Marco, R. Sevilla, Y. Zhang, J. J. Ródenas, and M. Tur, *Exact 3D boundary representation in finite element analysis based on Cartesian grids independent of the geometry*, International Journal for Numerical Methods in Engineering, 103 (2015), pp. 445–468.

[35] A. Menk and S. P. A. Bordas, *A robust preconditioning technique for the extended finite element method*, International Journal for Numerical Methods in Engineering, 85 (2011), pp. 1609–1632.

[36] T. S. Newman and H. Yi, *A survey of the marching cubes algorithm*, Computers & Graphics, 30 (2006), pp. 854–879.

[37] J. Nitsche, *Über ein Variationsprinzip zur Lösung von Dirichlet-Problemen bei Verwendung von*

Teilräumen, die keinen Randbedingungen unterworfen sind, Abhandlungen aus dem Mathematischen Seminar der Universität Hamburg, 36 (1971), pp. 9–15.

[38] J. Parvizian, A. Düster, and E. Rank, *Finite cell method: h- and p-extension for embedded domain problems in solid mechanics*, Computational Mechanics, 41 (2007), pp. 121–133.

[39] C. Pechstein and C. R. Dohrmann, *A unified framework for adaptive BDDC*, Tech. Rep. RICAM-Report 2016-20, Johann Radon Institute for Computational and Applied Mathematics (RICAM), Austrian Academy of Sciences, Altenberger Str. 69, 4040 Linz, Austria, 2016.

[40] J. Rudi, A. C. I. Malossi, T. Isaac, G. Stadler, M. Gurnis, P. W. J. Staar, Y. Ineichen, C. Bekas, A. Curioni, and O. Ghattas, *An Extreme-scale Implicit Solver for Complex PDEs: Highly Heterogeneous Flow in Earth's Mantle*, in Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, SC '15, New York, NY, USA, 2015, ACM, pp. 5:1–5:12.

[41] Y. Saad, *Iterative Methods for Sparse Linear Systems*, Other Titles in Applied Mathematics, Society for Industrial and Applied Mathematics, Jan. 2003.

[42] D. Schillinger and M. Ruess, *The Finite Cell Method: A Review in the Context of Higher-Order Structural Analysis of CAD and Image-Based Geometric Models*, Archives of Computational Methods in Engineering, 22 (2014), pp. 391–455.

[43] B. Sousedík, J. Šístek, and J. Mandel, *Adaptive-Multilevel BDDC and its parallel implementation*, arXiv:1301.0191, (2013).

[44] Y. Sudhakar, J. P. Moitinho de Almeida, and W. A. Wall, *An accurate, robust, and easy-to-implement method for integration over arbitrary polyhedra: Application to embedded interface methods*, Journal of Computational Physics, 273 (2014), pp. 393–415.

[45] A. Toselli and O. Widlund, *Domain Decomposition Methods*, Springer, 1 ed., Nov. 2004.

[46] T. Toulorge, C. Geuzaine, J.-F. Remacle, and J. Lambrechts, *Robust untangling of curvilinear meshes*, Journal of Computational Physics, 254 (2013), pp. 8–26.